Numerical Weather Prediction

Precipitation diagnostics for a high resolution forecasting system.

This is the stage 3 report from the Storm Scale Modelling project.
The project is half funded by Defra.

Forecasting Research Technical Report No. 423
Joint Centre for Mesoscale Meteorology Report No. 145

Nigel Roberts

email: nwp_publications@metoffice.com

©Crown Copyright

PROJECT DOCUMENTATION

PRODUCT DESCRIPTIONS

Stage 3 report from the Storm-Scale numerical modelling project: Precipitation diagnostics for a high resolution modelling system

Release: Draft
Date: 6th September 2003

PRINCE2

Author: Nigel Roberts

Owner: Nigel Roberts

Client:

Document Number: Version 2.0

2 Product Descriptions History

2.1 Document Location

This document is only valid on the day it was printed.
The source of the document can be found at:
\\Meto04\nwp_workgrps\Forecasting Research \JCMM\MesoscaleModelling\DEFRA_Storm_Scale_NWP

2.2 Revision History

Date of this revision:
Date of Next revision:

Revision date	Previous revision date	Summary of Changes	Changes marked
	22/07/03	First issue, version 1.0	
06/10/03		Ammended following comments by Brian Golding and Peter Clark, version 2.0	

2.3 Approvals

This document requires the following approvals. Signed approval forms are filed in the Quality section of the PCB.

Name	Signature	Title	Date of Issue	Version
B.Golding	*Brian Golding*			2

2.4 Distribution

This document has been distributed to

Name	Title	Date of Issue	Version

Contents

2 Product Descriptions History 2
2.1 Document Location 2
2.2 Revision History 2
2.3 Approvals 3
2.4 Distribution 3
3 Overview 5
4 Introduction 5
5 Why is post processing of high-resolution precipitation forecasts required 6
5.1 High resolution model factors 6
5.1.1 Limitations of a discrete gridlength 6
5.1.2 Larger-scale errors through the boundaries 7
5.1.3 Stochastic noise 7
5.2 Other factors 8
5.2.1 Raw output can look 'too believable' 8
5.2.2 Sensible methods for verification are required 8
5.2.3 User requirements 9
6 Precipitation products from existing nowcasting systems used in the Met Office 10
6.1 Nimrod 10
6.2 Convection Diagnosis Project (CDP) 11
6.3 Gandolf (Generating Advanced Nowcasts for the Deployment of Operational Land-based Flood forecasts) 13
7 Requirements and assumptions 14
7.1 Requirements 14
7.2 Assumptions about the model that can be utilised 15
8 Examples of products 17
8.1 Point probabilities 17
8.1.1 The proportion of gridpoints within a defined square area that have a maximum precipitation rate over a specified threshold. 17
8.1.2 The proportion of gridpoints within a defined square area that have a precipitation accumulation over a specified threshold. 21
8.2 Diagnostics on squares 28
8.2.1 Maximum precipitation rate within squares 28
8.2.2 Maximum precipitation accumulation within squares 29
8.3 Products for regions or catchments 34
8.3.1 Catchment average precipitation accumulations 34
8.3.2 Catchment maximum average precipitation accumulations 35
8.3.3 Probability of a catchment average accumulation exceeding a threshold 40
8.4 Graphical representation of catchment products 42
8.4.1 43
9 Final comments 44
10 References 45
11 Acknowledgements 45

3 Overview

The current development of high resolution numerical modelling should mean that in future it will be possible to forecast precipitation events with significantly greater accuracy on local scales than is currently achieved. The Met Office is now working on a 'storm scale' model with a horizontal gridlength of 1 km for forecasts up to 12 hours ahead. Although the available computer power is not sufficient at present for such a model to be run operationally, the time will come when it will be possible. To make best use of this advance in forecasting capability, there is a need to develop appropriate products that are designed to meet customers' requirements and expectations and also address the particular problems that arise in the interpretation of high-resolution model output. This report will discuss why there is a need to develop post-processing methods for the presentation of high-resolution precipitation forecasts, particularly for flood prediction and warning, and give some examples of what these products might look like.

4 Introduction

The Storm Scale Numerical Modelling project has now reached the stage 3 - to report on possible diagnostic outputs from the model runs that could be used to convey the information in a form suitable for use in flood prediction and warning.

The objective set out for the project as a whole was to investigate the ability of a storm scale configuration of the Met Office NWP model to predict flood-producing rainfall up to 12-hours ahead and to develop appropriate tools for interpreting and presenting the predictions so that they enhance operational flood prediction capabilities. In stage 1 it was demonstrated, using four case studies, that a 1, 2 or 4 km gridlength model is capable of producing realistic simulations of convective events which, for those particular cases, gave an improvement on the 12 km gridlength forecasts (Stage 1 report, Roberts 2003a). In general, the highest resolution simulations produced the best results. In stage 2, the sensitivities to changes in model parameters were tested (Stage 2 report, Roberts 2003b) during the development of the High Resolution Trial Model (HRTM) (Lean 2003) and a 'best at present' configuration found for a 1, 4 and 12 km gridlength nested system which can now form the basis for future studies. Further HRTM case studies have continued to show encouraging results (Lean 2003). It is now an appropriate time to consider the ways in which the output from high resolution forecasts should be interpreted and presented.

This report will firstly discuss (section 4) why there is a need to develop precipitation diagnostic products from high resolution forecasts (~4, 1 km gridlength) in the first place, when in the past it has been generally sufficient to use the fields that came directly out of the coarser resolution models (~60, 12 km gridlength). Partly this is to do with the characteristics of the numerical model and partly it is to do with the objectives of a storm scale modelling system. Secondly, in section 5, precipitation diagnostics that are already produced within operational Met Office nowcasting systems 'Nimrod' and the 'Convection Diagnosis Project' (CDP) will be described because the same ideas can be applied. Thirdly, in section 6, the requirements for appropriate products will be discussed. Fourthly, in section 7, examples of specific diagnostics that have been coded and could be used will be shown. This will be a subset of examples compared with the large number of possibilities that exist. It is not meant to be exhaustive or to define the precise diagnostics that should be used. The intention is to provide a starting point and ideas for discussion. Finally, section 8 will conclude.

5 Why is post processing of high-resolution precipitation forecasts required

Why bother to mess around with the precipitation fields that are output from a storm-scale model forecast system? Presumably, as forecasts become more accurate and resolve more features, the direct output from the model (which solves all the necessary equations) will become increasingly useful and sufficient? In fact, this is not necessarily the case and care needs to be taken in the interpretation of high-resolution-model output. The general problem is outlined below and then more specific reasons for requiring high-resolution precipitation diagnostics are described under two headings. The first are to do with the nature of high-resolution models, the second is to do with the purpose and interpretation of the output.

It is not possible to represent the exact state of the atmosphere at the start of a forecast. Any initial fields can only be a best estimate with an associated error. If forecasts are run from almost identical initial states, they will diverge and the differences can be explained in terms of a forecast 'error' or 'uncertainty'. Zhang et al demonstrate this in 3.3km gridlength simulations of convection and cyclogenesis. Since the initial state (including its error) is never completely known, we cannot regard any single forecast as the most likely possibility; it is only one realisation from an infinite number of alternatives with different likelihoods that form a probability distribution. This might give the impression that a single 'deterministic' forecast is pointless (some would argue that), but experience has shown that single forecasts are useful and this is partly because it is believed that different scales diverge at different rates and in so doing exhibit different characteristics.
The smallest scales diverge the most rapidly and eventually the error on these scales becomes the same for all forecasts regardless of the initial state. It can be thought of as random noise, which can be dealt with in a statistical way. This will be mentioned again under the heading of stochastic noise.
The largest scales grow the least during a forecast and are therefore the most predictable, thus making a single forecast useful for those scales. However, it is possible that the error on the larger scales could be big to start with if the initial state of the model is not very close to reality.
The intermediate scales have errors that can also grow rapidly and may cause the greatest difficulty in a high-resolution forecast because it is these scales that we would most like to be able to predict. We know we can't predict the smallest scales, except statistically and we don't need a high-resolution model for the largest scales.
In considering how to post-process high-resolution precipitation we should consider how to deal with two scales. (1) The scale over which we average to get a probability density function. Then we can account for the small-scale noise. (2) The scale over which we expect well-resolved features to be in error. We want to find the smallest scale that is generally predictable for a particular forecast system. Unfortunately, this will lie somewhere in the range of intermediate scales discussed above and will vary from event to event.

5.1 High resolution model factors

5.1.1 Limitations of a discrete gridlength

Any Numerical Weather Prediction (NWP) model is attempting to represent the continuous atmosphere on a discrete grid. It is clearly not able to resolve meteorological features that are smaller than a grid square, but as well as that, it is not capable of properly representing features that are smaller than around 5 grid squares across. Although not

specifically a problem for high-resolution grids alone (at first sight it may appear to be less of a problem at high resolution), there is an important difference in the way convective precipitation is generated in say a 1 km gridlength model compared to 12 km – and this is crucial because convection is a primary concern in a 'storm-scale' forecast system. On a 12-km grid, showery precipitation is mostly unresolved by the model dynamics (certainly individual storm cells are unresolved), and a convection scheme is employed to represent these smaller scales. Precipitation generated by a convection scheme is an average of scales that provide a large number of convective cells within a gridsquare and is only indirectly related to the model gridlength. On a 1-km grid spacing however, most of the convection is produced by the model dynamics and will not be properly represented if smaller than several gridsquares across. The implication of this is that precipitation output from a 1-km gridlength model forecast should always be averaged to a 5-km grid spacing, or larger, before being presented. The convective precipitation output from a coarse resolution model does not need to be averaged because the convection scheme has effectively done this already. The proviso here is that, in practise, convection schemes can also produce spurious gridscale structure that needs to be interpreted as such.

5.1.2 Larger-scale errors through the boundaries

The accuracy of a high-resolution model forecast will be constrained by the accuracy of the information that comes through its boundaries. It may be for example, that a forecast will develop a realistic band of thunderstorms some 50 km away from where they occurred in reality simply because the dynamical ascent and convective instability that was required for the storms to trigger came through the edge of the domain misplaced by ~50 km. The error arises in the coarser resolution forecast that supplies the boundary conditions, but can have a serious impact on the usefulness of a storm-scale forecast. A displacement error of say 50km for a rain area in a 12km gridlength model is only around 4 gridlengths and is therefore expected to occur at times.

There is a need then, to take account of this uncertainty in some way. It is probably the most difficult unknown to take account of using post processing of the output, but should nevertheless be attempted. Ideally, a very large high-resolution domain should be used to help reduce boundary contamination, but in practise, computational cost forces a compromise. A useful approach might be to run an ensemble of the coarser resolution forecast to help determine the uncertainty – but that is outside the remit of this particular piece of work.

5.1.3 Stochastic noise

As the resolution of numerical models improves, more and more of the smaller meteorological features (e.g. small showers) can be resolved, but these smaller scales are inherently less predictable and more transient. Small changes to a model's parameters, land surface fields or initial conditions will change a forecast on these scales (and perhaps larger scales too). Experiments with the convection scheme, model diffusion and microphysics (Roberts 2003, Forbes and Clark 2003) have already shown this to be true. The envelope region of convective activity may be predictable (assuming a reasonable initial state and boundary information), but the individual showers are not, and could be thought of as random noise if viewed over a large enough area. For this reason, it is not sensible to extract the detail in high resolution model output at face value: a forecast that says it will rain at a particular location 3 hours from now, but be dry in a village 2-km away cannot be believed to that level of accuracy. This is true whatever the resolution of the model (and is different to the problem outlined in 4.1.1).

It is not true to say however, that showers exhibit completely random behaviour: there might still be collective characteristics that are well within the capacity of a numerical model to

predict. The intensity of the rain, the likely size and duration of showers, their spacing, the adjacency to orographic features, the most active period and the region(s) most at risk are all aspects of a forecast that should be more predictable than the individual rain elements themselves. So it might be valid to say, for example, that it is 'very likely there will be showers over the north of a river catchment in 2-3 hours time and that they are expected to be heavy and lead to a high risk of localised flash floods'. That is how a forecaster might express such a situation and it is this kind of information which should be conveyed by suitable precipitation diagnostics that take account of stochastic noise and yet exploit the more predictable behaviour.

5.2 Other factors

5.2.1 Raw output can look 'too believable'

The trouble with a high-resolution model precipitation forecast is that it can look very much like a radar picture. In one way this is a good thing because it is, after all, what we would like a model to be able to achieve. The problem is that anyone viewing the output could be very tempted to believe the small-scale detail in a forecast - and we know from the discussion above that this is not a wise thing to do. It would therefore, be dangerous to provide raw output to customers who are not experienced with high-resolution forecasts because the information could easily be misinterpreted. We need a variety of ways of presenting precipitation forecasts that not only meet user requirements, but are also tailored to their level of expertise.

5.2.2 Sensible methods for verification are required

The main conclusion from section 4.1 is that high-resolution model output is very likely to be misleading if simply taken at face value. For the same reasons, any attempt to verify the performance of a high-resolution model forecast on spatial or temporal scales that are inherently unpredictable will be meaningless.

If we wish to assess the accuracy of high-resolution model forecasts (and we do), it is not appropriate to adopt the traditional verification method of comparing forecasts against single-point observations at specific times. The small-scale random nature of precipitation or even very small displacement errors in well resolved features can make a forecast score poorly, when in fact, there might be a great deal of useful and accurate information. It should also be said that it is not ideal to verify coarser resolution forecasts in this way either because of the representivity problem of comparing point values with large grid-square averages.

An alternative way to evaluate high-resolution model precipitation forecasts is to compare precipitation diagnostic products with radar data that has been processed in the same way (or, if a large number of events the radar data does not require the same processing). This has several benefits. (1) The products should have been designed to extract the believable information and deal with the uncertainty in a forecast so the verification is worth doing. (2) If the products are designed for customers to use, then the verification is dealing with the product the customer sees, so any changes in performance will have a direct significance to the customer. (3) The verification can be used in the development and tuning of the products themselves – this is an important feedback. (4) Radar gives the best spatial and temporal view of 'reality' and although there can be significant errors, these errors too could be accounted for in the processing.

This report will not say much more about verification issues, but it is important and will crop up again in future Storm Scale Modelling project reports.

5.2.3 User requirements

Different users/customers/software systems will extract from a precipitation forecast the particular aspects that are most applicable to them. A hydrologist might need to know about the amount of rain that will fall over a river catchment because of a serious flood risk, whereas others might be more concerned about whether there will be a dangerous deterioration in motorway driving conditions, or whether it is likely to rain this evening or not, or if rain will save England from test match defeat. To best meet the needs of users/customers it makes sense to generate products that are tailored to their requirements (within the model accuracy constraints already mentioned). If this is not done, we will end up with a high-resolution forecasting system that misses its potential as a tool for practical decision making.

6 Precipitation products from existing nowcasting systems used in the Met Office

Before presenting potential diagnostic precipitation products for high-resolution NWP, it is worth having a look at what is produced from existing nowcasting systems to see if the same ideas can be applied.

6.1 Nimrod

Nimrod is a nowcasting system that merges extrapolation of precipitation derived from radar and satellite imagery with mesoscale model forecasts (gridlength ~12km) up to 6 hours ahead. A large number of products are routinely generated by the operational system. The ones to do with rainfall are described below.

1. Straightforward output of rainfall rates on the Nimrod grid (5x5km gridsquares) every 15 minutes during a 6-hour forecast.

2. Output of hourly precipitation accumulations on the Nimrod grid for a 6-hour forecast.

3. Probability of rain occurring at each hour during a 6-hour forecast - A square is defined around each gridpoint of size determined by a derived wind error (calculated elsewhere in the Nimrod system). Then the fraction of gridpoints within each square that are precipitating is calculated to produce a probability of rain occurring. The probabilities are taken at single snapshot times every hour from T+1 to T+6. The squares increase in size with forecast time. See example in Figure 1(a)

4. Probability of rainfall rates over 0.5 and 4.0 mm/hour at each hour during a 6-hour forecast - Identical to 3. - but calculating fractions of gridpoints with precipitation rates greater than 0.5 or 4 mm/hour.

5. Maximum rainfall accumulations within a square surrounding each point within a 1-hour period during a 6-hour forecast - The hourly rainfall accumulation at each point is displayed over a square of size determined by the wind error. Squares will overlap; in which case the highest value is displayed. See example in Figure 1(b)

6. Catchment statistics for previous 3 and 24 hour precipitation accumulations.

The 3 and 24 hour accumulations over each 5x5km gridsquare are displayed over catchment areas and statistics presented along with a warning status.

Example of text information:

Region: East Midlands, Time: 15:30 30/07/03
Accumulation period: 3 hours
Threshold: 15mm
Catchment Max: 25.31mm
Catchment Average 2.27mm
Pixels exceeding threshold: 8
Pixels within 75%: 10
Warning Status: 2

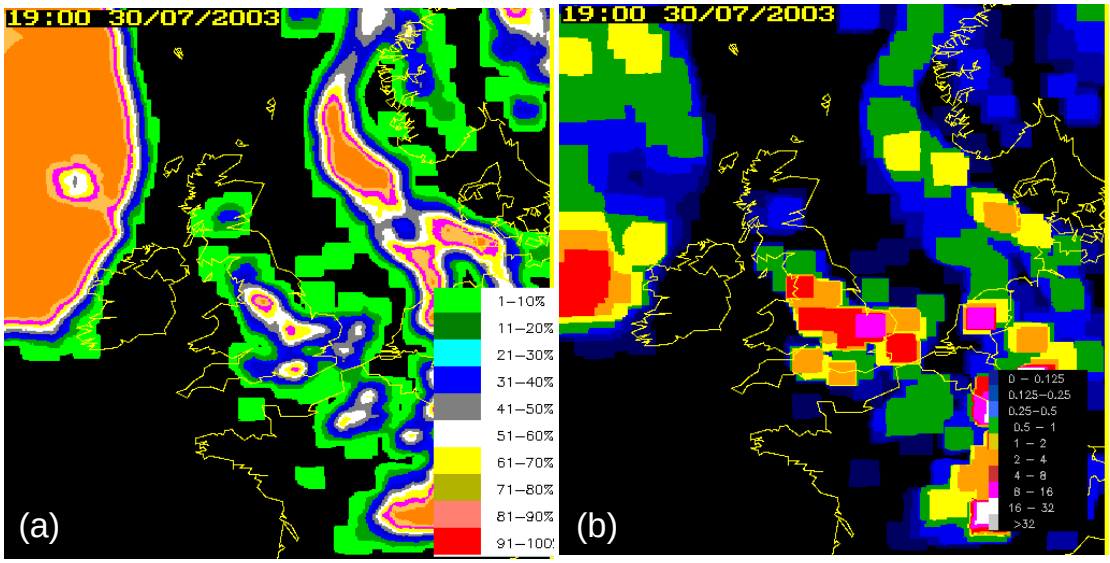

Figure 1. Examples of precipitation products from Nimrod. Information is given in the text

6.2 Convection Diagnosis Project (CDP)

The CDP is a system that uses mesoscale model forecast fields (gridlength 12km) to determine the chances of convective clouds developing given expected model errors in temperature, humidity and dynamical forcing. A convective-cloud life-cycle model is then used to predict the development, structure and movement of any showers. A range of products has been developed and the operational ones are described below.

1. Peak rainfall rates expected in showers

Model gridsquares are coloured according to the largest rainfall rate that could occur somewhere within a shower at its present stage of development. Low values will be indicative of weak convection and areas of high values will indicate regions of deep or strong convective activity. It provides an indication of the "flashiness" of the convective precipitation. Locally heavy downpours can seriously effect driving conditions. See example in Figure 2(a)

2. Probability of getting a shower with rates >0, >2, and >10 mm/h at any one location chosen at random inside each NMC Press region sometime during the 6 hour period prior to the verification time.

The UK is divided into 30 Press Regions and the regions are coloured in according to the probability of getting a shower at any one spot chosen at random within that region and time period. The purpose is to present the chance of any one spot getting at least one shower during the period. The product takes into account all shower movement and development during the 6-hour period. See example in Figure 2(b)

3. Maximum rainfall accumulation per shower in each NMC Press Region during the 6-hour period prior to the verification time for probabilities of individual shower occurrence >= 10% and >= 70%.

This product displays the maximum amount of rain that could be collected at any one spot for each shower that occurs. Large, long-lasting showers will produce relatively large accumulations, whereas fast-moving showers will only give low accumulations. It is a useful product for identifying Regions at risk from heavy, slow moving showers. Probability thresholds are taken from product 2. A larger probability of occurrence will be associated with fewer showers, as fewer showers will be expected to that degree of certainty. See example in Figure 2(c)

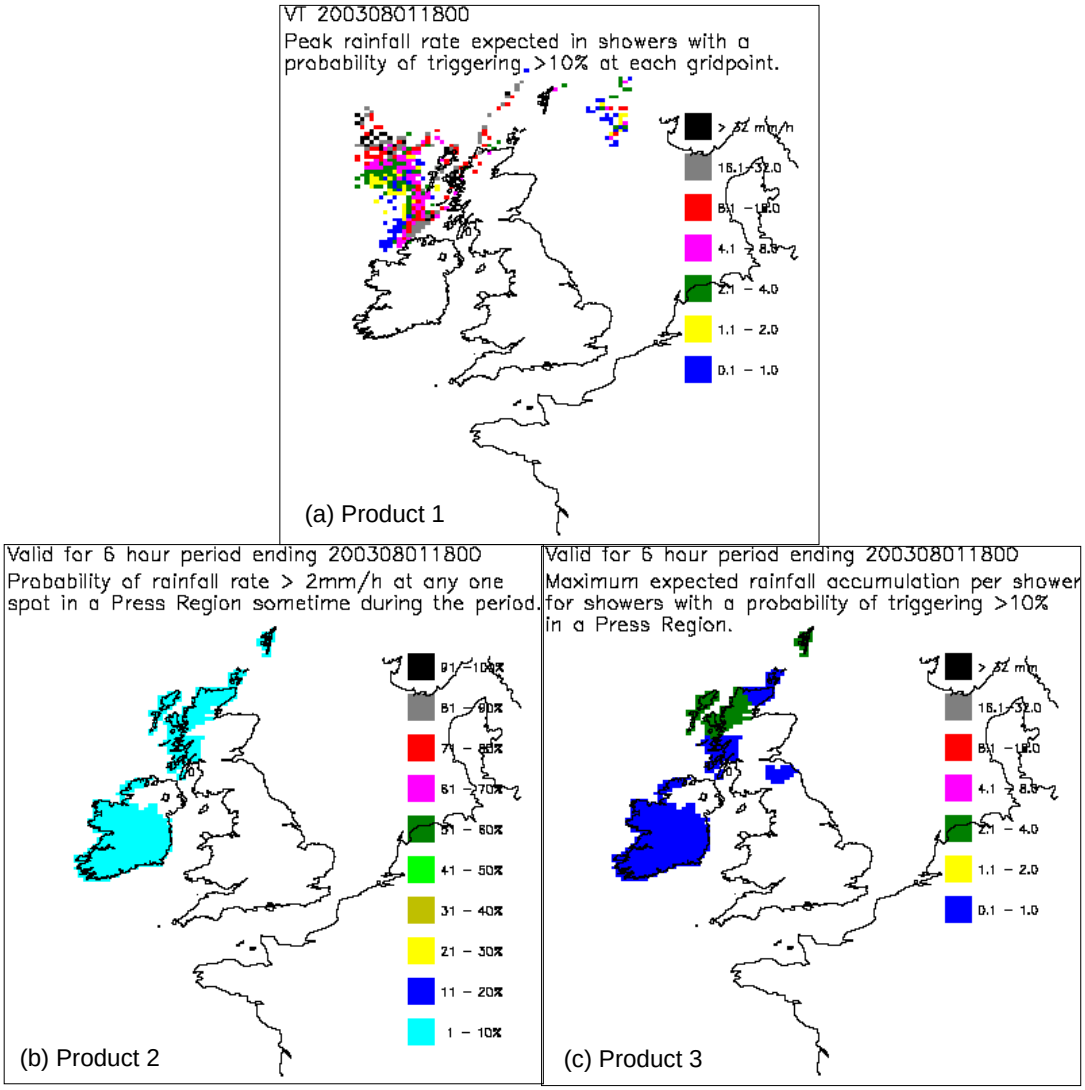

Figure 2. Examples of products from the Convection Diagnosis Project (CDP). Information is given in the text.

6.3 Gandolf (Generating Advanced Nowcasts for the Deployment of Operational Land-based Flood forecasts)

Gandolf, like Nimrod, is also a nowcasting system that merges extrapolation of precipitation derived from radar and satellite imagery with mesoscale model forecasts. It was developed in partnership with the Environment Agency (EA) and operates on a gridlength of 2km rather than 5km used by Nimrod. Other differences with Nimrod are the use of a neural network cloud classifier, an object orientated cloud model for convective precipitation (like the CDP) and an improved advection scheme.

Products include instantaneous rain rate, and 15 minute rain accumulation forecasts with a range of three hours. In addition, 5 minute, plain text warnings of heavy rainfall within specified EA regions or river sub-catchments are automatically generated. At present, no probabilistic products are produced. However, a new version of Gandolf has been developed that will run an ensemble of nowcasts and this will be used to produce probability distributions of rainfall scenarios.

7 Requirements and assumptions

This section will outline some of the requirements for high-resolution-model precipitation products and the assumptions that are needed in developing products described in the next section.

7.1 Requirements

1. Retain and enhance only believable information on scales of interest

One of the primary aims of precipitation products should be to enhance the characteristics of a high-resolution forecast that are thought to be most believable and at the same time remove the scales that are too small to be reliable. This might be, for example, to include information about precipitation rates but smear out the spatial distribution of the rain areas.

2. Quick to calculate

If these products are to be used in a real time situation, then they must easy to calculate quickly. This is particularly true if a storm-scale modelling system is set up to produce frequent forecasts for only a few hours ahead. Such a system would then have to impose a restriction on the number of precipitation diagnostics that are routinely produced.

3. Comprehensible to users

It is no use developing something that is so contrived nobody understands what it means. The aim of developing these products is to have something that is of practical use, rather than as a purely academic exercise. However, to be over simplistic is not helpful either – the real world is complicated and many users have a high level of understanding. A range of products will be needed to convey the information – again, not so many that the time to produce them becomes prohibitive or that information overload causes confusion.

4. Used in verification

As outlined in section 4.2.2, the performance of any new storm-scale modelling system must be carefully examined. Diagnostic products need to be incorporated into this process and should be an integral part of the comparison of model performance against observations or other models and nowcasting techniques. The verification process should also be used to examine the usefulness of the products themselves.

5. Useful to users

Products must be useful; otherwise there is no point in producing them. Having said that, it impossible at this stage to say precisely how useful any particular diagnostic tool might be. Dialogue with users, in addition to objective verification scores will determine this eventually. For now, it is enough to consider the potential value to users of new products and state what that is anticipated to be.

6. Scientifically sound

There is a danger of developing diagnostic tools that in some way try do the job of the model itself. The aim is not to make forecasts more accurate. Products should be designed to present the most useful aspects of a forecast in ways that take account of expected model errors. It is really just making better use of the output rather than changing it. For example, it would be tempting, but incorrect, to move a predicted area of rain west by 20 km because the model had often incorrectly generated the rain 20 km too far east in a particular region. Even if such an error was seen in many previous situations it does not necessarily mean that it will re-occur on other occasions in the same way. An error of that kind is either symptomatic of a problem with the formulation of the model, inaccuracy in the initial conditions or to do with the representation of features in the model like hills or urban areas – and it is those problems that would need addressing. It is however, valid and sensible to develop products that are designed to account for the likelihood of ~20 km displacement errors in predicted rainfall patterns.

There will be other more subtle ways in which it is possible to inadvertently generate questionable products.

7.2 Assumptions about the model that can be utilised

Most of the methods of precipitation diagnostics to be discussed will depend to some extent on the validity of two assumptions about high-resolution model behaviour.

1. The correct type of precipitation structures are developed

Most of the precipitation diagnostics are supposed to take some account of possible displacement errors in model forecasts. It is only reasonable to do this because we believe that the model usually simulates the right kind of weather event even if not necessarily positioned correctly. In frontal situations the model will nearly always generate frontal rain rather than say scattered showers or an isolated thunderstorm. Many years of experience with numerical models have shown this to be true at coarser resolution – otherwise there would not have been the success there has been. At high resolution, the simulations that have been performed so far, display the same characteristics. This is not surprising, since the larger-scale meteorological environment is inherited from the coarser model, but even on the smaller-scales, showers tend to organise into bands or be small and scattered when they should be.

If we can largely depend on the model to get the meteorological scenario correct, then we can develop diagnostic products that deal with errors in size, location, timing and occurrence of precipitation features. The characteristics of these errors can be understood. If the model frequently got the meteorological scenario wrong, it would be too unreliable to make the production of diagnostic products a worthwhile exercise. Post-processing would not be useful in the face of such a fundamental problem with the model.

2. Rainfall rates are reasonably realistic

If the model generally gets the correct type of rainfall features, it might be reasonable to assume that the precipitation rates are sufficiently realistic. By distinguishing correctly between, say, thunderstorms or frontal bands, the model can generate either intense or moderate rain appropriately. Not only that, the model is controlled by physical constraints that limit the possible errors in rainfall rates. Indeed, case studies have shown that the rainfall rates from the model are usually of the right magnitude when compared

with radar, but that there are also some problems. The rates can be very inaccurate locally when the model develops single 'grid point storms' with extreme rainfall. This is more serious in a 4km-gridlength model than a 1 km, but ongoing modifications to model formulation and tuneable parameters should reduce this problem dramatically and also improve the representation of precipitation generally.
If gross errors in precipitation intensity can be largely removed then it is valid to say that errors in precipitation rates are much less significant than errors in precipitation location and products can be developed on this basis. Once it is assumed that the model will generally produce sufficiently realistic rainfall rates, the same is probably true for rainfall accumulations, provided that the accumulation period is long enough to take account of small-scale spatial variability.

8 Examples of products

This section will explain some products that have been produced so far and show examples. The examples will make use of high-resolution model simulations of two rainfall events that have already been described in the first Storm Scale Modelling project report (Roberts 2003a). They are (1) A quasi-stationary band of convective precipitation that persisted over Southeast England for over 18 hours on 11-12th October 2000 and produced severe flooding in parts of Sussex and Kent. Accumulations measured by some rain gauges exceeded 100mm. (2) Scattered showers and thunderstorms on 3rd May 2002. These showers were generally small, but did produce some locally very heavy rain.

No mention will be made of averaging high-resolution model output to a coarser resolution grid; not because it isn't worthwhile doing, but because an explanation of the method is not needed.

8.1 Point probabilities

Probabilities can conveniently represent the uncertainty in a precipitation forecast and greatly add to the information that a forecast can deliver. However, it is recognised that the use of probabilities is sometimes unpopular because it is not always clear how best to use them. This can hopefully be overcome with greater familiarity. Provided that probabilities are not presented in complete isolation and can be interpreted for decision making, they should be an invaluable part of a high-resolution-modelling post-processing system.

8.1.1 The proportion of gridpoints within a defined square area that have a maximum precipitation rate over a specified threshold.

Method

Over a specified time period, sampling at defined intervals, the maximum precipitation rate that occurred at each gridpoint is found. A square of size n*n gridpoints is defined. For every gridpoint, the n*n square is centered over that gridpoint and the fraction of surrounding gridpoints within that square which have a maximum precipitation rate over a specified threshold is found. The fractions can be interpreted as the probability of the precipitation rate exceeding a given threshold over each model gridsquare within the time period. See Figure 3

Purpose

The purpose is to try and determine the probability of intense precipitation events that may, for example, cause flash floods, make driving hazardous, ruin important sporting/public events or have some other adverse effect on society/industry.

Sampling within a square around each point is meant to take account of spatial errors and grid-scale representivity errors in a forecast. It is assumed that the precipitation rates are sufficiently reliable, and at gridlengths of 1 km (or less), they are representative of the rates a rain gauge might measure. On a 4km grid however, rainfall rates from the model will be averages over the larger squares and probably not representative of rain-gauge measurements – particularly in convective situations. This diagnostic is probably not

appropriate for models with gridlengths longer than 2km unless some kind of additional statistical processing is done to determine corresponding rainfall rates on a 1km grid.

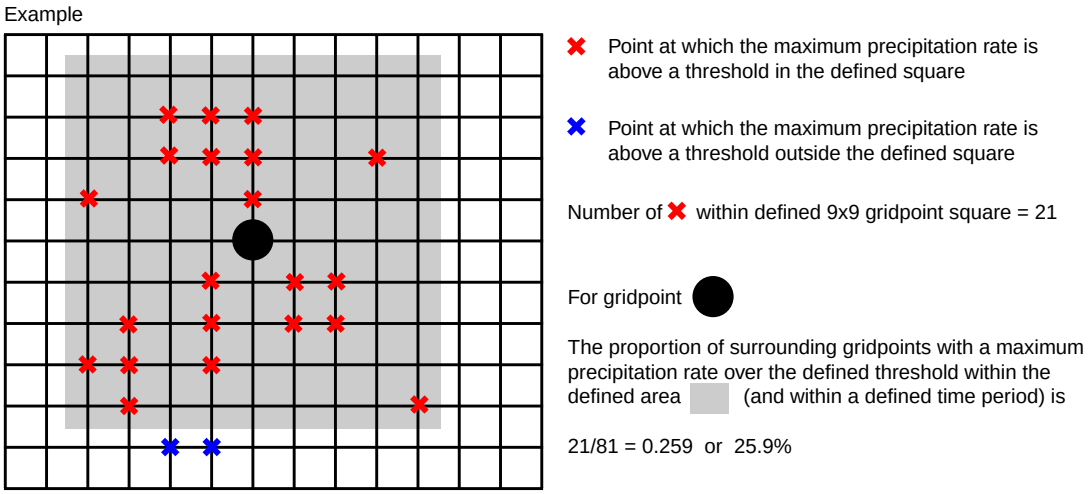

Figure 3. In the above example, if the gridlength is 4km, then the large grey square is covering an area of 36x36km. The chance of finding a gridpoint with a maximum precipitation rate over the specified threshold, in that 36x36km square, is 25.9%. That is the value assigned to the gridpoint marked with the black circle. A value for every gridpoint is calculated in the same way.

Parameters

1. Size of the square

The size of the square should be determined by the forecast skill for each sized square, the resolution of the model, the forecast lead time, the resolution of the model that provides the boundaries, the data assimilation system that provides the initial conditions and very importantly the purpose of the output. Forecast skill can be determined by comparing with radar data over a large number of events. It is expected that skill will improve with square size, but this has to be balanced by the decreasing value of the information as the square size increases. The size of the squares is a trade off between greater precision from small squares and greater reliability from large ones. There is probably no point in going to the expense of running a 1-km gridlength forecast model and then generating probabilities derived over 100x100km squares.

For a 4km gridlength forecast out to T+24->36 driven from a 12km gridlength model. The square should be at least 36x36km because this is 3x the 12km model gridlength. It should probably be larger than 48x48km to take better account of the errors in the 12km model, analysis errors and to get better statistics from a larger number of points in the square. However, if very large squares are used, the information from a 4km-gridlength run may be too 'washed out'. E.g. a square covering the whole England would have good statistics but be irrelevant for Yorkshire.

For a 1km gridlength forecast out to T+6->12 driven from the 4km gridlength model. It is assumed that the start of a forecast would be more accurate for small features and the short forecast time means that the accuracy of the 4km-boundary model is less important. The square in this case could be much smaller, but at least 12x12km.

2. Time period

A sensible time interval could be anything between 1-12 hours because it will contain a large enough number of sampling times. For a 1km gridlength run 1-6 hours seems more appropriate. A fast-moving feature travelling at 50 km/hour will have travelled 300km in 6 hours and in doing so could cross a 1-km gridlength model domain. The time period used should also be determined by the purpose of the output.

3. Threshold

The threshold used depends on the user requirements and also the gridlength of the model. For a 1km-gridlength forecast, thresholds of 10, 20, 40 and 80mm/hour may be appropriate. At 4km gridlength the model will not be able to resolve the rainfall intensities of the 1km run for small-scale features and so lower thresholds that have been shown to be statistically comparable may be used – or else this product may not be deemed appropriate for 4km.

Example

An example of generating probabilities of rain rates exceeding 16mm/hour using 40x40km squares for a 1-km gridlength forecast is shown in Figure 4(a). This is the scattered convection event that occurred on 3rd May 2002.

Extensions

1. Circles

Instead of squares, it makes more sense to use circles, so that a uniform maximum sampling distance from each central point is maintained.

Example

Figure 4(b) shows the effect of using circles with a radius of 40km. The output looks smoother and 'nicer' at the edges of low-probability regions, but otherwise the impact is not significant except that the magnitude of the higher probabilities has increased. This is because, in this example, the number of sampled points has decreased.

2. Soft threshold

Instead of stating that gridpoint precipitation rates must be at or over the threshold to be counted as a hit in the probability calculation, it is reasonable to think that if a rainfall rate is just below the threshold it could still be counted but given less weight i.e. a value less than 1. Any model value of precipitation rate will have an error associated with a level of inherent unpredictability in a forecast. Given a slight change in model parameters or initial conditions, a gridpoint value might change from being just below the threshold in one forecast to just above in another. The magnitude of the precipitation rates is just as much part of the model uncertainty as the spatial distribution of rainfall – albeit considered less of a factor.

Example

Figure 4(c) shows probabilities from circles the same size as in Figure 4(b) but gridpoints with precipitation rates less than the threshold but greater than 2/3rds of the threshold are counted as a fraction of a hit according to a linear transition between 0 for 2/3rds threshold and 1 for the threshold or above. So, for the threshold of 16mm/hour, 2/3rds of the threshold is 10.67mm/hour and if a gridpoint had a precipitation rate of 13mm/hour

it would be assigned a value of (13-10.67)/(16-10.67) = 0.44. This function is somewhat arbitrary, but reasonable, given that we don't know what the real error function looks like. The main effect in this example (and in general) is to increase the probabilities mostly in the regions where they were already comparatively higher. An alternative would be to have a symmetric soft threshold in which a function also applies to rates above the threshold.

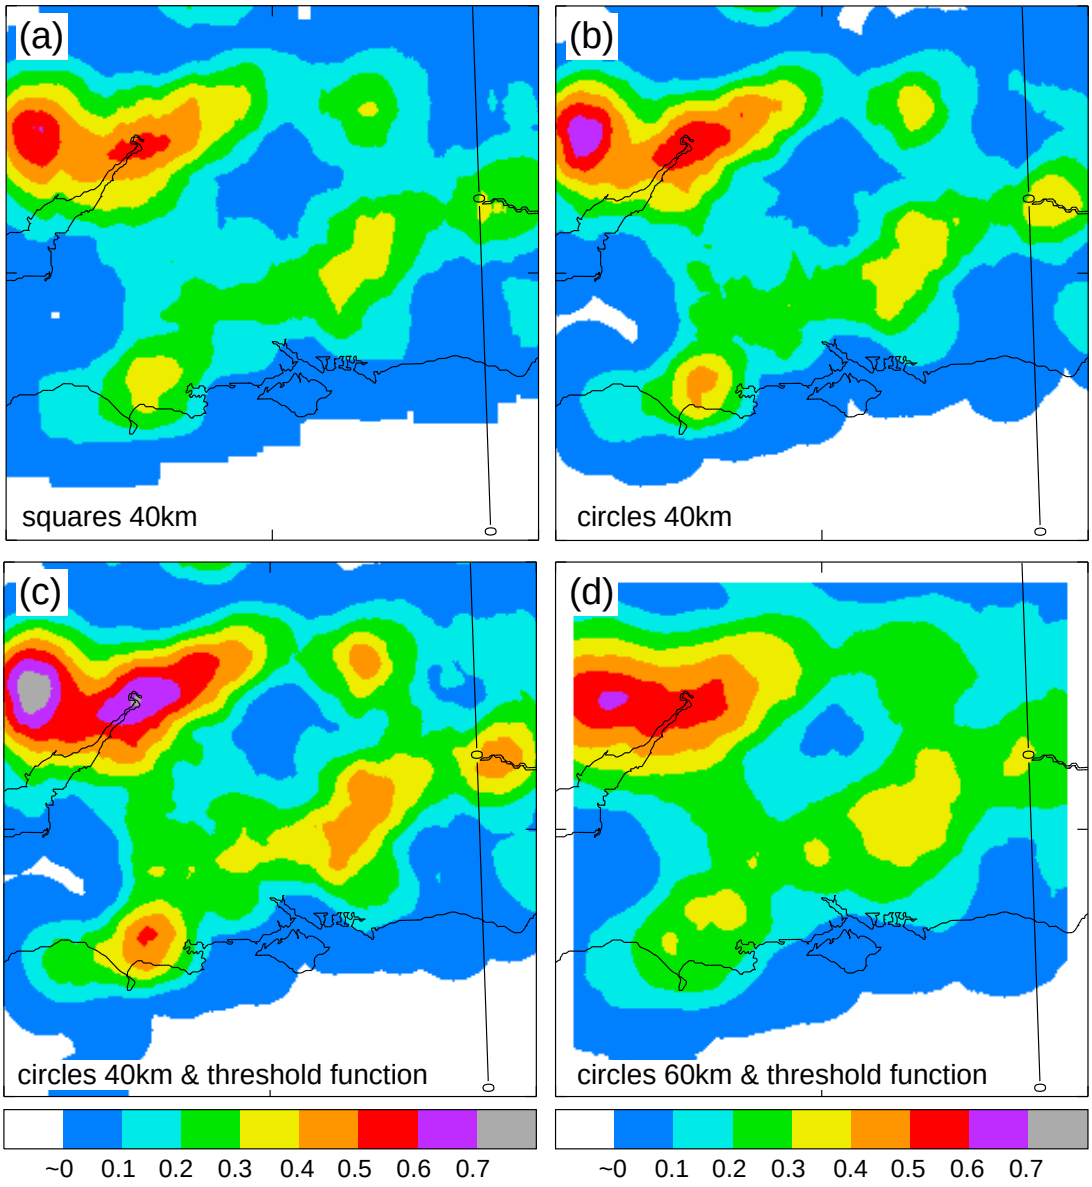

Figure 4. Examples of probability products from a 1km gridlength forecast starting at 01UTC 03/05/02 for the 9-hour period 09 to 18UTC. Probabilities of precipitation rates exceeding a threshold of 16mm/hour. Calculations used (a) squares of length 40km, (b) circles of diameter 40km, (c) circles of diameter 40km with a function for the threshold, (d) as (b) but with 60km diameter circles. See text and Figure 3 for more information.

8.1.2 The proportion of gridpoints within a defined square area that have a precipitation accumulation over a specified threshold.

Method

This exactly the same as 7.2.1, except that precipitation accumulations are used instead of maximum precipitation rates.

Purpose

The purpose is to try and determine the probability of high rainfall accumulations that may lead to flooding.

Parameters

The size of the square or circle and the time period - The same as for probabilities of precipitation rates.

Threshold

As for probabilities of precipitation rates, the threshold used depends on the user requirements and also the gridlength of the model. Unlike precipitation rates, the threshold for accumulations should not depend on the model gridlength to the same extent, given a long enough sampling period.

Further extensions

These apply to both accumulations and rates

1. Combine different forecasts

The probabilities from forecast can be combined with the probabilities from a previous forecast (or other ensemble members) to give a better indication of uncertainty in the forecast – particularly due to errors coming through the boundaries. Weighting can be applied to extract more information from the most recent forecast – or a smaller square or circle can be used for the most recent forecast.

Example

Figure 5 shows how probabilities from two forecasts can be combined so that more account of the true uncertainty can be taken.

2. Convert from probabilities to odds

Probabilities can be converted to odds by: Log(odds) = log(p/1-p). The advantages of doing this are (1) Odds are probably better understood by the general public than probabilities. (2) The distribution becomes Gaussian from minus infinity to plus infinity and therefore adds detail to the very low or high probability regions. (3) It is better for some forms of verification. See examples in Figure 5(d) and Figure 8(b)

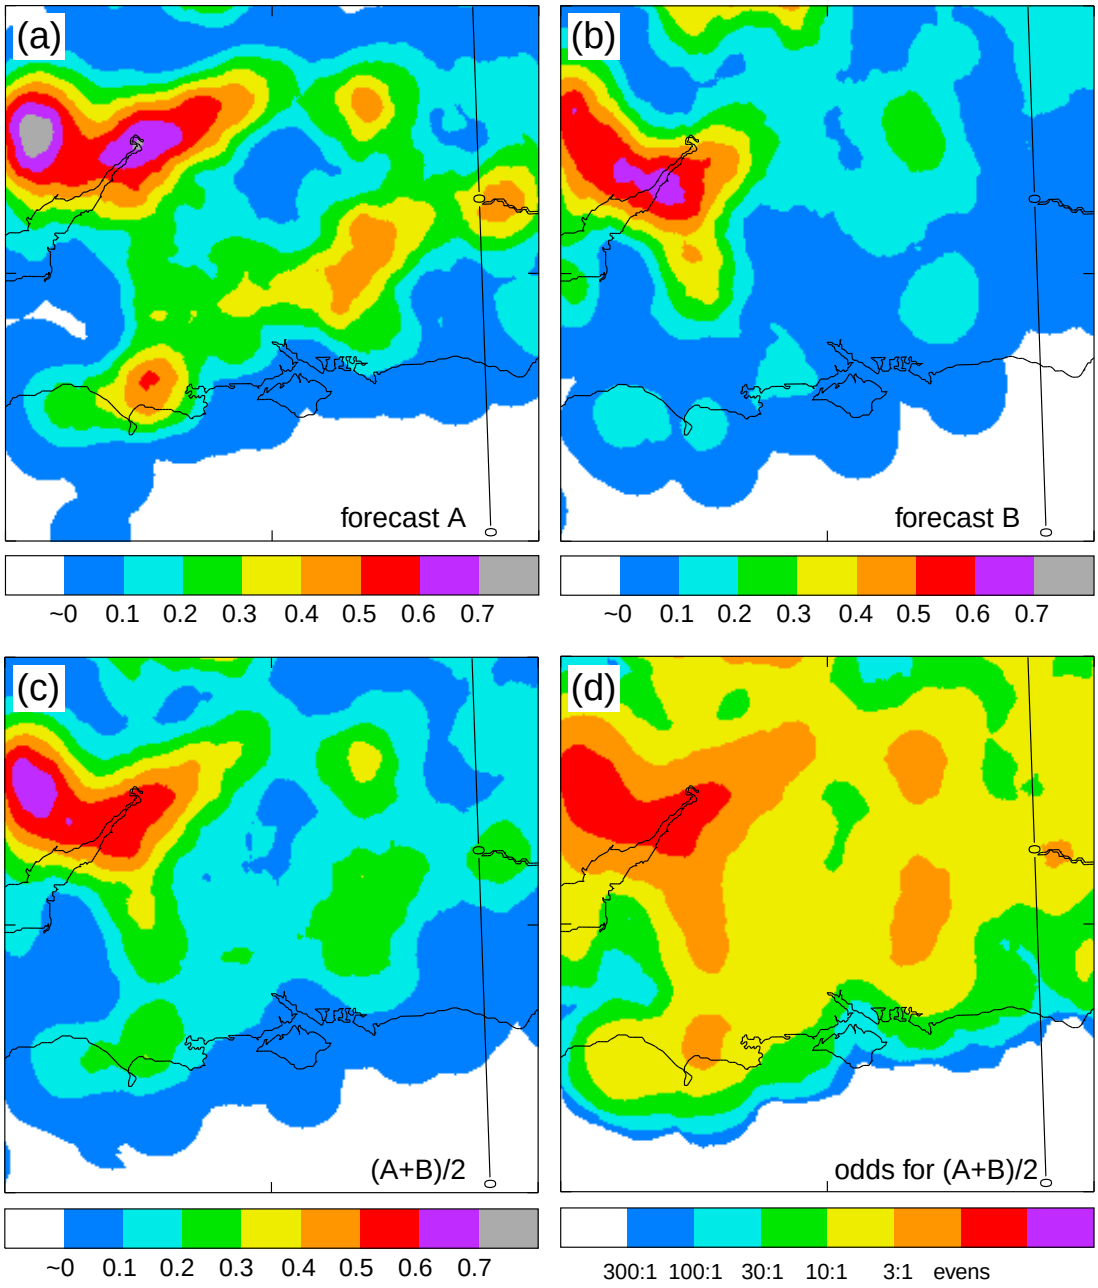

Figure 5. Examples of probability products from two 1km gridlength forecasts starting at 01UTC 03/05/02 for the 9-hour period 09 to 18UTC. (a) Probabilities of rainfall rates exceeding 16mm/hour, computed using circles of diameter 40km and a soft threshold (see text) from forecast A; (b) As for (a) but from forecast B which had different information coming through the boundaries; (c) a combination of (a) and (b); (d) probabilities in (c) converted into odds. Forecasts A and B differed in that they had information coming through their boundaries from two different 4-km gridlength forecasts. The run supplying the boundaries to forecast A had less resolved convection than the one supplying forecast B

3. Combined squares or circles

An advantage of using squares or circles is that it is simple, and especially for squares, computationally efficient. A disadvantage is that there is no tail off in the weight that is assigned to each point at some distance from the centre of the square or circle. The result is,

for example, that a point 20km away from the centre of a 20km radius circle may be just as important as one 1km away, but a point 21km away is not taken into account at all. One way of producing a tail at the edge is to combine several squares or circles in the way shown in Figure 6(a) & (b). The probabilities are computed at every point for several different sized squares or circles and then combined by averaging over the number of square sizes used. Figure 6(b) shows a schematic example of the rectangular probability distributions obtained for several sized squares, given the 10% hits within the inner square and none outside that is depicted in Figure 6(a). The blue line in Figure 6(b) shows the probability distribution from the average of the square distributions, with a flat constant probability within the smallest square and a tail either side.

The main problem with this technique is computational expense. It is clumsy to repeat the probability calculations many times over increasingly large and therefore computationally more expensive squares or circles when the application of a weighting function over the largest square or circle would suffice. For this reason, it may not be ideal as a 'real time' product for customers. The benefit might be in verification. An important part of evaluating high-resolution model performance will be to compare probabilities generated over different sized squares or circles with radar processed in the same way to see how the forecast skill varies with square or circle size. This will mean that the probabilities are produced anyway and can be easily and quickly combined for further verification.

4. Use of a recursive filter

Another way of generating probabilities is to use a recursive filter. The aim is to 'dress' each point so that it represents the centroid of a plausible distribution. The recursive filter provides a cheap way to do this.

Method

1.Each gridpoint with a value over the specified threshold is given a value of 1.0 (a hit) and every other gridpoint a value of 0.0 (a miss).

2.The recursive filter is then applied in the positive x direction, then the negative x direction, then the positive y direction then the negative y direction. This may be repeated. The way the recursive filter works is shown in Figure 7.

Parameters

(1) The weighting parameter alpha - The larger closer this value is to 1, the more the information is spread from gridpoint to gridpoint and the more the smaller scales are removed.
(2) The number of passes - The more passes, the more information is spread from gridpoint to gridpoint and the more the smaller scales are removed. Both parameters could be tuned to represent a certain amount of spatial filtering in the same way that different sized squares can be used in the other method.

The boundaries

The value at the boundary should be determined by what is outside the boundary as well as what is within the domain and this is unknown. The edge could be dealt with by either applying a function to compute what the value would be if there were a large number of points with no hits beyond the domain, or by using a buffer of extra points with no hits around the domain. A more realistic way might be to provide information from outside the domain with data interpolated from the coarser resolution model that provided the

boundaries for the high-resolution forecast. This is appropriate for precipitation accumulations, and should take the problem away from the edge of the high-resolution domain, but is less applicable for precipitation rates because the coarser resolution forecast will be unable to represent the higher rainfall intensities.

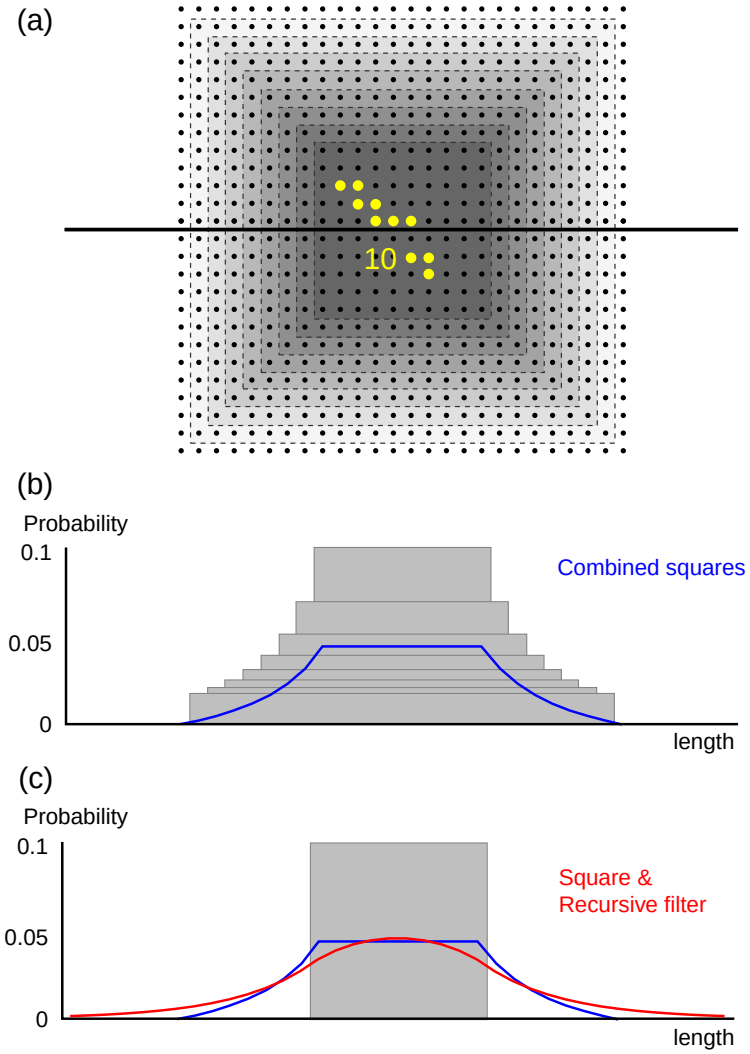

Figure 6. Schematic diagram to show how a combination of probabilities computed from different sized squares can be used to produce a probability distribution with a fall off from the edge of the small square to the edge of the large square. (a) shows, as an example, 10 points that have been counted as hits on the grid (yellow points) and 8 squares (dashed lines), increasing in size, placed around those points. (b) shows the probability distributions (grey rectangles) that are generated along the line in (a) from calculating the proportion of hits (yellow points in (a)) for every square. The blue line shows the combined probabilities when the probability at every point, from every square, is added together and divided by the number of squares. (c) shows the probability distribution from the smallest square in (a) (grey rectangle), the blue line is the same as in (b), the red line shows the probability distribution computed when starting with the grey rectangle and applying a recursive filter once in the x, -x, y and –y directions with parameter α=0.8. See text and Figure 7 for more information.

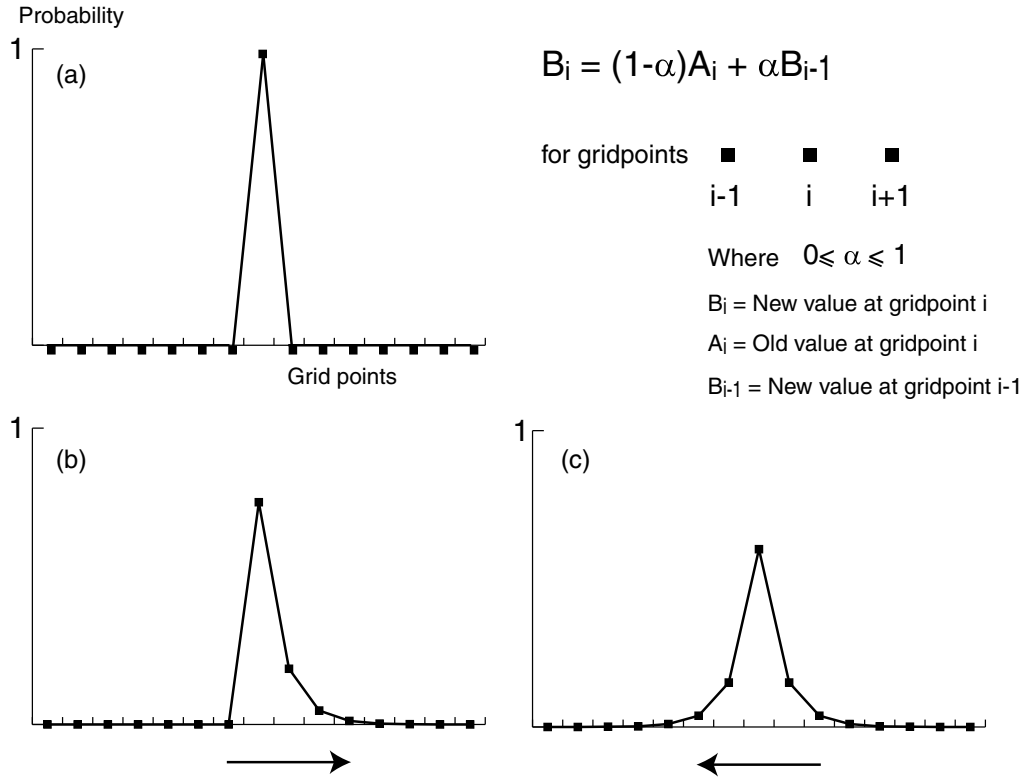

Figure 7. The graphs show the effect of applying the recursive filter to a single gridpoint value of 1 in just the x direction with two forward and backward passes. The first graph (a) shows the initial state. The next two graphs (b) & (c) show the effect of applying the recursive filter in the forward then backward x direction.

A combination of a forward then backward pass in the x direction followed by a forward then backward pass in the y direction will give an isotropic spreading out of the data. This can be used as a way of producing a probability distribution of where a variable is likely to exceed a threshold. If a single gridpoint has a probability of 1 and all others 0 to start, then this can be smoothed to give a new probability of say 0.4 at the gridpoint itself, 0.2 at adjacent gridpoints tailing off towards zero some distance away.

5. A combination of the square (or circle) method with a recursive filter

The method of generating probabilities by finding fractions over squares or circles can be combined with the recursive filter. First generate the probabilities by using squares or circles of n*n gridpoints as described in 7.2.1, then apply a recursive filter. The outcome is to produce a probability pattern which is like that obtained with the combined squares (or circles) method, but with a longer spread (Figure 6(c)). It is computationally efficient and only requires one pass in each direction of the filter because the square sampling method has already smoothed out the smaller scales. This approach is regarded as the best one out of those described here for generating probability maps.

Examples

Figure 8(a) displays the probabilities from the combined circles method for the 03/05/02 case with circles increasing in size from 40km diameter to 60km diameter. A comparison with Figure 4(c) & (d) shows that this approach has maintained much of the information from the 40km circles as well as the spread from the 60km circles. Figure 8(b) shows what the probabilities look like when converted into odds.

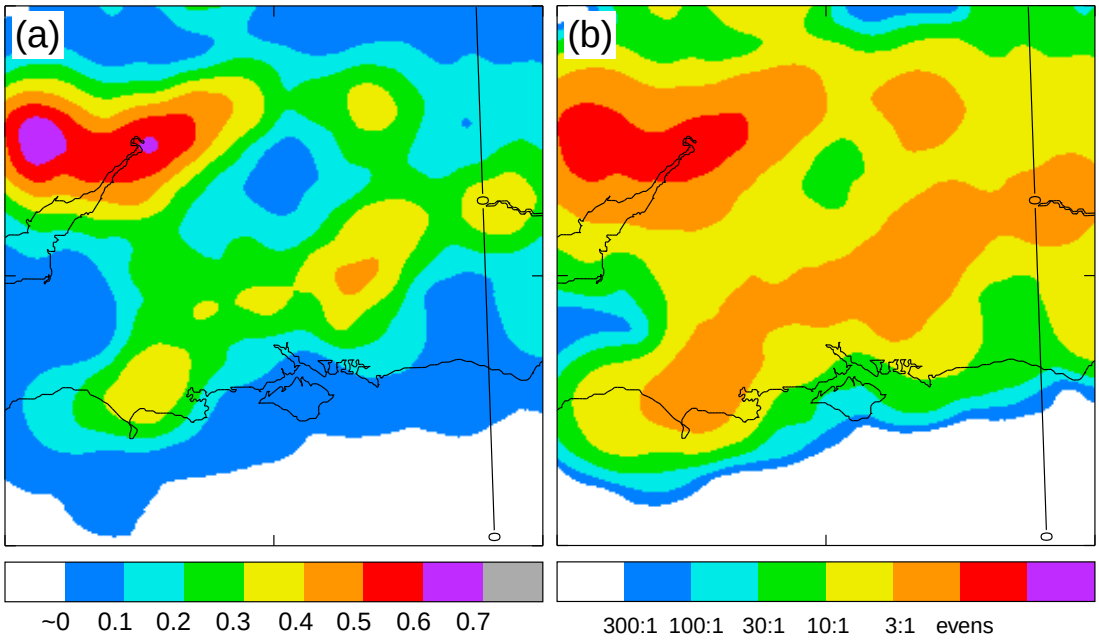

Figure 8. Examples of products from a 1km gridlength forecast starting at 01UTC 03/05/02 for the 9-hour period 09 to 18UTC. (a) Probabilities of rainfall rates exceeding 16mm/hour, computed using circles of diameter 40km to 60km and a soft threshold (see text); (b) Odds for rainfall rates exceeding 16mm/hour from probabilities P in (a). Colour scale represents log(p/(1-p)).

Figure 9 shows the use of probabilities for rainfall accumulations. The example is from a 2km-gridlength-model simulation of the persistent rainband that gave severe flooding over parts of Essex and Kent on 11-12 October 2000. Figure 9(a) shows the probability of rainfall accumulations exceeding 30mm using 60km diameter circles and weighting given to accumulations >20mm. Figure 9(b) is the same except that 84km diameter circles have been used. The effect of using the larger circles is to increase the size of the region with a probability > 0, but also to reduce the highest probabilities from >0.75 to >0.6. Figure 9(c) shows what happens when circles with a diameter between 60 and 84km are combined. The region with probabilities >0 is the same as for the 84km diameter circles, but the maximum probabilities are now between those produced by the 60 and 84km circles. Figure 9(d) shows the effect of using a 60km diameter circle and then a recursive filter with one pass in each direction and alpha=0.8. This is very similar to Figure 9(c) except that the lowest probabilities have been spread out more and it was very much quicker to compute.

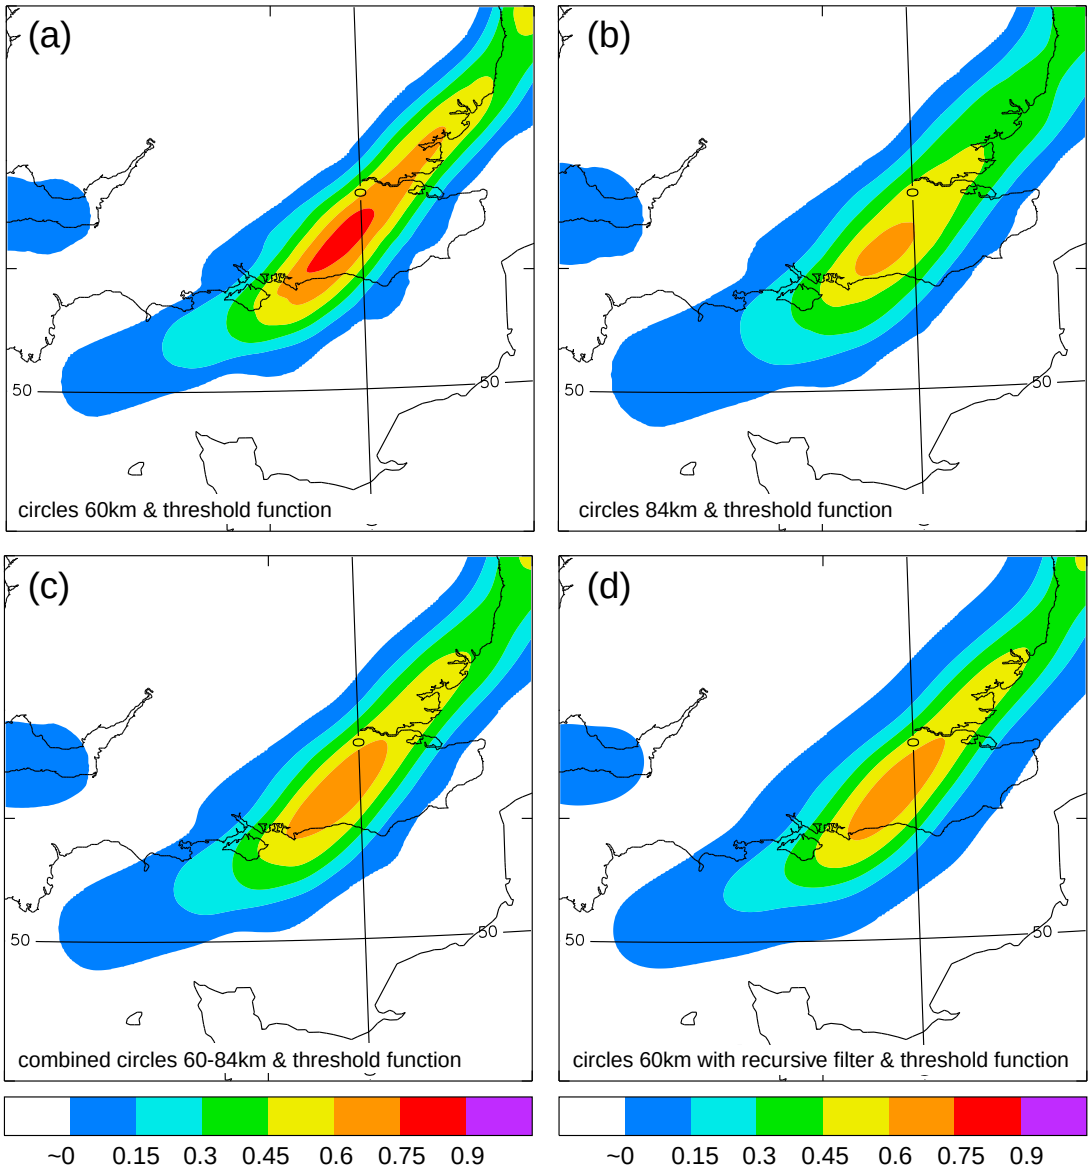

Figure 9. Examples of probability products from a 2-km gridlength forecast starting at 18UTC 11/10/00 for the 18-hour period 18UTC 11/10/00 to 12UTC 12/10/00. (a) Probabilities of rainfall rates exceeding 16mm/hour, computed using circles of diameter 60km and a soft threshold; (b) As for (a) but using 84km circles instead; (c) As for (a) but using a combination of probabilities from circles increasing in size from 60 to 84km; (d) As for (a) but using a recursive filter with a single pass in the x, -x, y and –y directions and α=0.8 applied to probabilities computed using circles of diameter 60km. See text for more information.

8.2 Diagnostics on squares

8.2.1 Maximum precipitation rate within squares

The maximum precipitation rate within a defined square area over a specified time period.

Method

Over a specified time period the maximum precipitation rate that occurred at each gridpoint is found (taken at short, preferably timestep intervals). A square of size n*n gridpoints is defined. These squares are tiled across the grid and the maximum precipitation rate that occurred within each n*n square is found. Where a square is not complete because of the edge of the grid, a null value is set. See Figure 10.

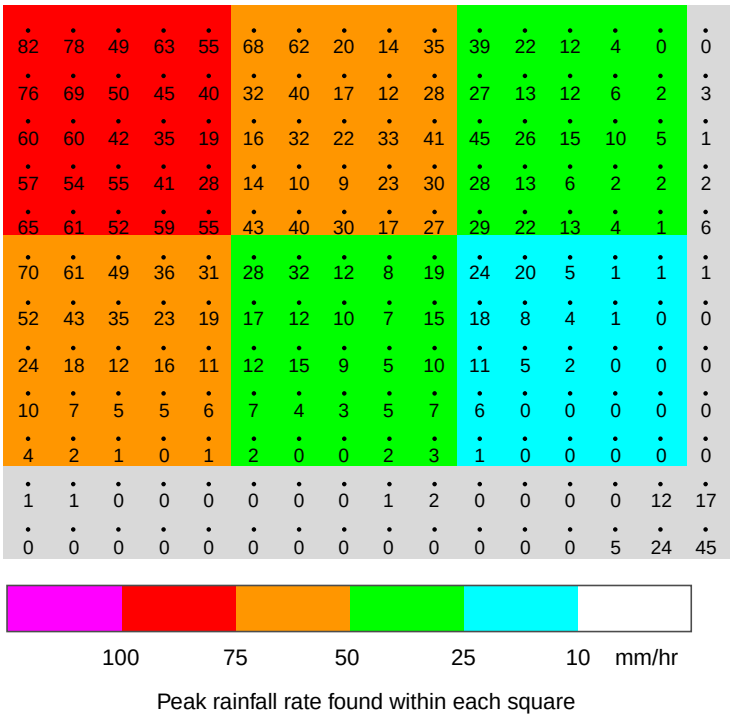

Figure 10. A schematic example of the maximum precipitation rates within squares. The red square indicates that at least one of the gridpoints within that square has a precipitation rate of between 75 and 100mm/hr. Three gridpoints with values of 82, 78 and 76mm/hr are within that range. The grey region shows where complete squares are not possible within this domain.

Purpose

The purpose is to give an indication of the peak precipitation rate that is possible at any given location within a square area over some time period. This is based on the assumption that high-resolution forecasts (~1-2km gridlength or less) produce precipitation rates that are comparable with point values that are recorded by rain gauges. It should be presented alongside probability information to give some indication of how likely certain precipitation rates are. It is very similar to the CDP product 1 mentioned in section 5.2.

Parameters

1. Size of the square

As for calculating probabilities, the size of the square should be determined by the forecast skill for each sized square, the resolution of the model, the forecast lead time, the resolution of the model that provides the boundaries, the data assimilation system that provides the initial conditions and the purpose of the output. It is probably best to match the size of square used here with that used in the probability products.

2. Time Period

Again, the choice of time period(s) will depend primarily on how the product is to be used and the size of the squares. The longest should be 12 hours and the shortest 1 hour. Any longer and it is not very useful, any shorter and we need to be sure that the model has that level of accuracy in terms of timing, location and triggering of convection.

8.2.2 Maximum precipitation accumulation within squares

The maximum precipitation accumulation within a defined square area over a specified time period.

Method

The same method as used to generate maximum precipitation rates within squares

Purpose

The purpose is to give an indication of the highest accumulation of precipitation that is possible at any given location within a defined area over some time period.

Parameters

As for maximum precipitation rates within squares

Example

Figure 12(a) shows an example from a high-resolution simulation of the 11-12/10/00 event. The squares with accumulations above 60mm immediately give an impression of a potentially serious event and the regions where this could occur.

Extensions

Extensions (1) and (2) could equally well apply to precipitation rates.

1. Use geographical areas instead of squares

In the same way that the CDP uses press regions, this product could be presented for press regions, or counties or river catchments. User requirements would determine whether this was worth doing.

2. Sample over larger squares

Instead of searching for the maximum accumulation within a square, look for the maximum accumulation within a larger square, but still present the information on the smaller square. The reason for doing this is to make the values in the squares more closely fit the positions of the point by point accumulations. This can be explained by looking Figure 11; the position of the highest accumulation is denoted by the X and this occurs outside the smaller square. If the sampling was just done over the smaller square, then point X would not be picked up even though the area of the smaller square is clearly at risk from that accumulation. If the larger sampling square is used, the accumulation value at x would be picked up and displayed over the smaller square (as well as the square to the right). An example of this method is shown in Figure 12(b) for squares that are the same size as the squares in Figure 12(a) but sampling over squares that are 4x the area. The picture can be changed considerably by altering the size of the larger sampling squares relative to the display squares. If the sampling squares become much larger than the display squares, the information is effectively displayed on the larger squares.

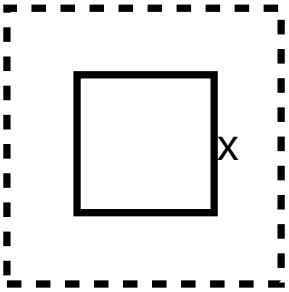

Figure 11. Schematic to show the use of a larger sampling square when finding the maximum precipitation rates or accumulations. See text.

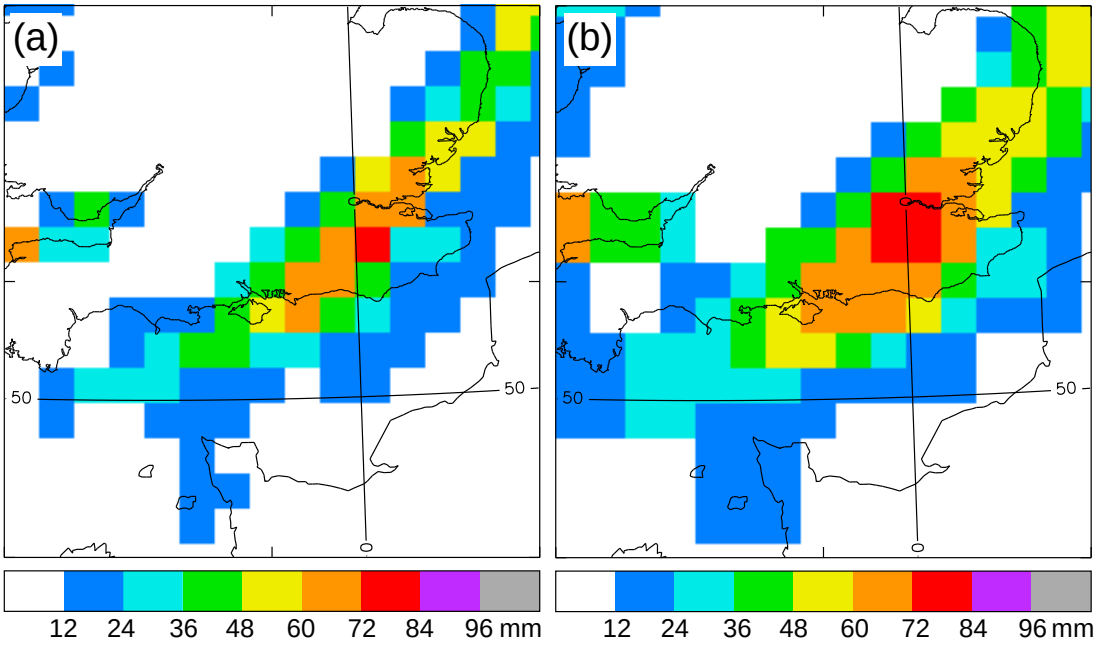

Figure 12. Examples of products showing maximum accumulations within squares from a 2-km gridlength forecast starting at 18UTC 11/10/03 for the 18-hour period 18UTC 11/10/00 to 12UTC 12/10/00. (a) Maximum accumulations found within squares (b) Maximum accumulations when sampling over larger squares. See text.

3. Obtain maximum extreme accumulations

In some convective situations, locally very high rainfall accumulations can occur because of a single very slow moving shower, or quasi-stationary band of showers or a succession of heavy showers. This may well be captured by a high-resolution model – especially if tied to orography in some way. However, because of the unlikely nature of extreme events, it is probable that exact realisations will be missed even though the situations that could produce them should be predictable. A way of creating an extreme event from the model output is to make use of a cell-locating algorithm to get the highest possible accumulations at each point and then present this over square areas.

Method

1. For a given gridpoint sampling over a time interval (usually a model timestep) find the precipitation rate.
2. Use a cell-locating algorithm to find nearby points within the same cell. The cell could be defined as comprising connecting points with a precipitation rate of over 3/4 the initial gridpoint value and above a pre-defined minimum value. Nearby could be defined as up to three or four gridpoints away.
3. Find the maximum precipitation rate within the nearby cell and assign that to the initial gridpoint. Figure 13 shows a schematic example of the effect of this re-assignment
4. Do steps 1 to 3 for all gridpoints and over all times.
5. For every gridpoint, add together the maximum precipitation rates at that gridpoint from every time and convert to accumulations
6. Locate the maximum accumulations within square areas and display

					1	15	
1	9					2	
7	3			10	5		
	5	1	13	4	9	8	
			2				3
		1	4				
						2	
				1	6		3

					1	**15**	
1	**9**					(15)	
(9)	(9)			(13)	(13)		
	5	*1*	**13**	(13)	(13)	(13)	
			(13)				(10)
		1	4				
						(6)	
				1	**6**		3

Figure 13. Schematic example showing how rainfall rates are re-assigned if the rates in each cell are substituted by the highest rate in that cell. The grid on the left shows rain rates before any changes and the grid on the right after re-assignment. The numbers indicate rainfall rates in mm/hour, a square with no number has no rain falling, a square can only be part of a cell if the rain rate is >1mm/hour. The numbers in circles show rates that have been changed because a higher value is within the same cell and less than 4 squares away. The bold numbers are maxima within cells. The numbers in Italics are unchanged because they are too small to be part of a cell.

Purpose

The aim is to see what the result would be if the highest precipitation rate within a rain area occurred everywhere nearby within that rain area and do the same for all showers over a period of time. The effect is to produce much higher rainfall accumulations in situations when the rainfall pattern is noisy i.e. convective and slow moving, which is when very high local accumulations are most likely to occur. It has some resemblance to the maximum accumulations found from each shower produced by the CDP.

Examples

Examples of the use of this product are shown in Figure 14(b) and Figure 15(b) as a comparison with the maximum accumulations displayed in Figure 14(a) and Figure 15(a). The values have been increased considerably in both cases and give a clear indication of a risk of very high totals. The values exceeding 150mm in the October 2000 event are not wildly off the mark, as 153mm was recorded, however, it is not possible to know whether totals > 60mm occurred in the May 2002 case. Verification of this product is extremely difficult because of insufficient numbers of observations and also because the aim is to give an indication of what might happen rather than predict what will happen. A large number of events would be required to determine how often predictions are realised and be able to tune the output. Despite this, it still has the potential to give a useful indication of the magnitude of locally high precipitation accumulations for customers who might want this information.

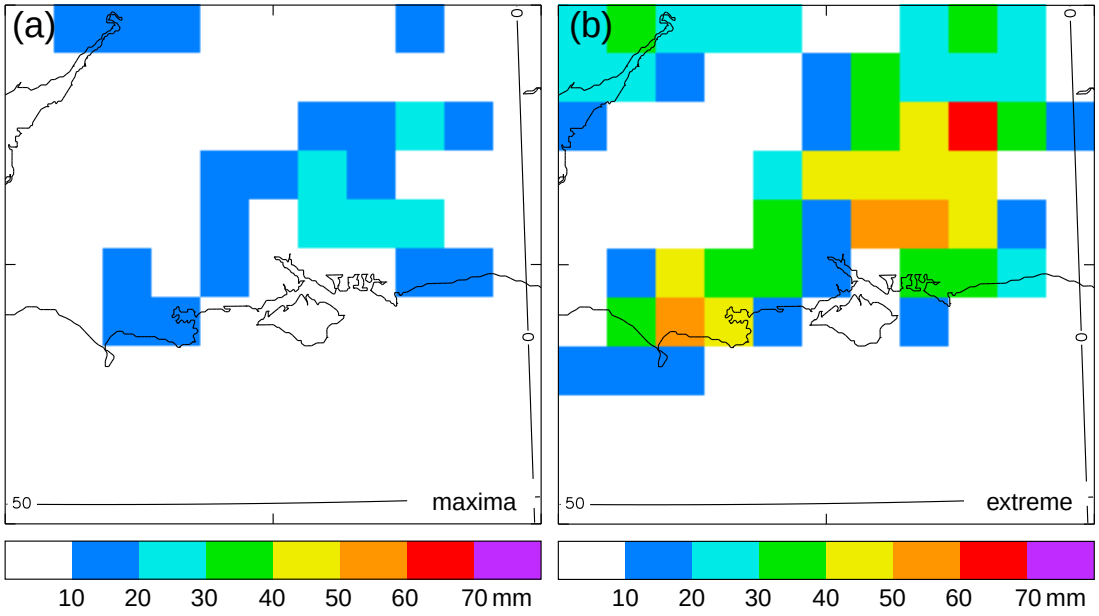

Figure 14. Examples of products showing maximum precipitation accumulations from a 1km gridlength forecast starting at 01UTC 03/05/02 for the 9-hour period 09 to 18UTC. (a) Maximum accumulations found within squares (b) Extreme maximum accumulations found within squares. See text and Figure 13 for more information.

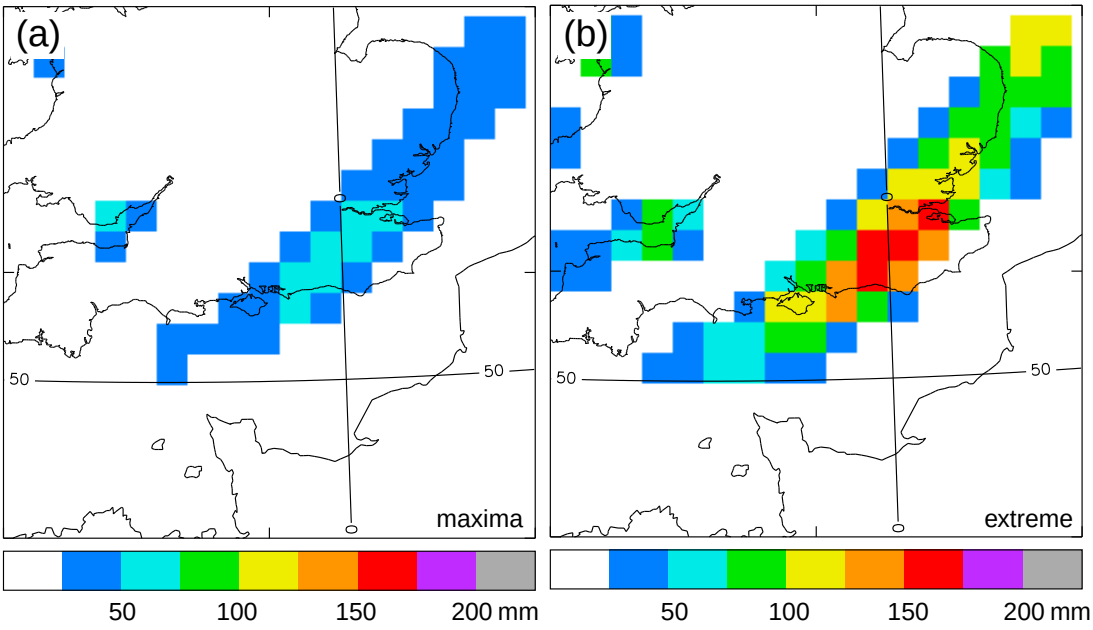

Figure 15. Examples of products showing maximum accumulations within squares from a 2-km gridlength forecast starting at 18UTC 11/10/03 for the 18-hour period 18UTC 11/10/00 to 12UTC 12/10/00. (a) Maximum accumulations found within squares (b) Extreme maximum accumulations found within squares. See text and Figure 13 for more information.

8.3 Products for regions or catchments

The ideas from the previous products can be applied to river catchment areas. Catchment products will now be described with examples from 2km gridlength forecasts of the 11-12th October 2000 event.

8.3.1 Catchment average precipitation accumulations

The most obvious product to generate is the average precipitation accumulation over each catchment.

Method

Find the total precipitation accumulation at each gridpoint over a time period. Add together the values at all points within a catchment and divide by the number of points in the catchment to get the mean. Do this for all catchments.

An example is shown in Figure 16(b). This is a useful product. It is quick to compute and can be verified against catchment average radar accumulations, but it does not give any indication of the spatial uncertainty in a forecast.

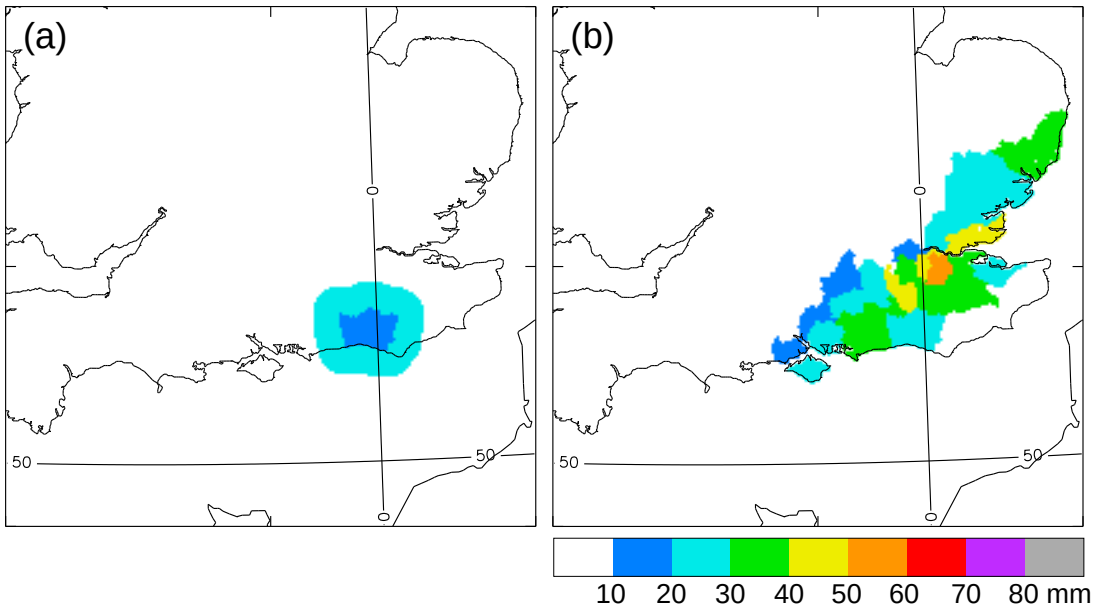

Figure 16. (a) A single river catchment area (dark blue) with a region 20km wide enclosing the catchment (light blue); (b) The average accumulations to fall over river catchments over the 18-hour period 18UTC 11/10/00 to 12UTC 12/10/00 from a 2-km gridlength forecast starting at 18UTC 11/10/00.

8.3.2 Catchment maximum average precipitation accumulations

There may be occasions when an area of high rainfall totals in a forecast will straddle several catchments. In such a situation the average accumulations may not be particularly high in any given catchment. However, if the forecast was slightly wrong and in reality the region of high totals was displaced, there might be considerably higher average accumulations in one or more catchments. The method below describes a way of generating a product that could warn of such a possibility.

Method

1. Find the precipitation accumulations over a specified time period at every gridpoint in the forecast domain
2. Locate the gridpoints in a catchment area and define a set of gridpoints surrounding the catchment that are within some pre-determined distance from the edge of the catchment. Figure 16(a) shows an example of a river catchment (dark blue) with an extended area of gridpoints that are within 20km of the edge of the catchment (light blue).
3. Count the number of gridpoints N within the catchment area alone.
4. Search through all the points in both the catchment area and extended area and find the highest N values of precipitation accumulation. Add these values together and divide by N to get the average of the maximum N accumulations.
5. Display this value on the catchment area map (or other form of presentation)
6. Repeat 1-4 for all catchments

Computation is quick.

Purpose

The aim is to present highest catchment average accumulations that could occur over a period; given that there will be some spatial errors in the precipitation forecast. This could be used as an indicator of when there is a possibility of catchment average accumulations reaching a critical level so that appropriate action can be taken.

Parameters

There are two tuneable parameters – (1) the length of the accumulation period and (2) the width of the extended area. The length of the accumulation period should be defined by user requirements. A reasonable width for the extended area is between 15 to 60km for a 1km-gridlength model. Less than 15km and the product is not accounting for plausible displacement errors. Greater than 60km and we are attributing a larger error to the model than it should usually have – and in doing so are more likely to get extremely high values that are very unlikely to occur.

Examples

Figure 17 shows two examples of applying this method for extended areas of 20 and 40km width. These should be compared with the catchment averages presented in Figure 16(b). The accumulations from this product are considerably larger than the straightforward catchment averages and become larger as the extended area becomes larger. The average accumulation for the catchment shown in Figure 16(a) has increased from >20mm in

18hours to potentially >40mm for a 20km extended area and >50mm for a 40km extended area.

One of the nice characteristics of this method is that the smaller catchments are more sensitive than the larger ones and this helps with a forecasting problem. The problem is that, the smaller the river catchment the more the accuracy of a catchment-scale forecast is sensitive to displacement or timing errors in the forecast. This is not helpful since it is the smaller catchments that have the shortest response times and require the most accurate forecasts. This method will at least provide a warning of the potential for a serious flood-producing event over small catchments even if the forecast misplaces areas of precipitation on the scale of the catchment.

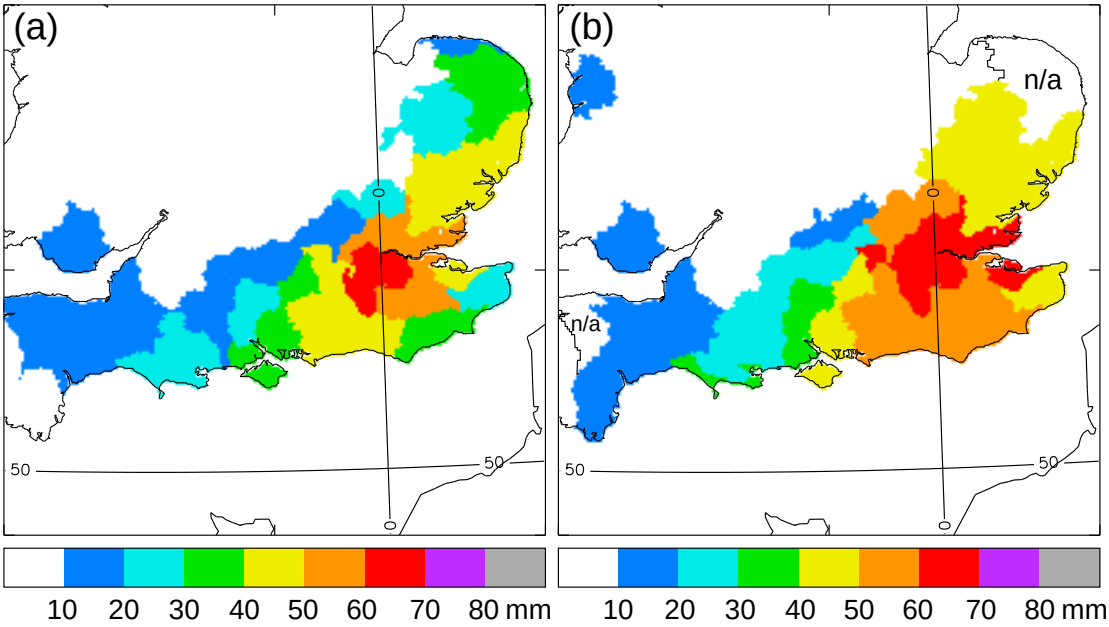

Figure 17. Maximum average catchment rainfall accumulations over an 18-hour period starting at 18UTC 11/10/00 from a 2-km gridlength forecast starting at 18UTC 11/10/00. Accumulations were computed using (a) a 20km extended area and (b) a 40km extended area. See text for more information.

Extensions

1. Take account of the distribution of the precipitation

There is a flaw in the calculation of the catchment average maximum accumulations. The method is supposed to take account of displacement or timing errors in a model forecast, but what it is really doing is lumping the highest accumulations from a large area into a smaller catchment area. That is fine if all the highest accumulations are mostly located together, then it amounts to translation of that region into the catchment area. However, if the highest accumulations are scattered about, the result is a convergence into the catchment area and that does not make sense as a realistic alternative scenario. Figure 18 will help to explain this. In the diagram on the left the predicted area of highest accumulations (shaded grey) is not entirely within the catchment area (solid square), but it is conceivable that given a slightly different forecast it could be located over the catchment area and give a much higher catchment average. In the diagram on the right, the highest accumulations also cover the same area as the catchment, because the method defines that they must, but are scattered over the extended area. In this situation, if the average over the grey squares is

taken as a possible average over the catchment then it is the same as squashing together all the grey squares into the catchment. In meteorological terms that might be the same as turning the accumulations from scattered showers into accumulations from one large single rain area. Since it has already been assumed that the model generally gets the meteorological scenario correct, this is inventing a rainfall pattern that should mostly not occur. That is not to say that a different realisation that produces a much higher catchment average accumulation is not possible. It is saying, that in this situation the occurrence of all the highest accumulations within the catchment is extremely unlikely. A method is needed to allow for different meteorological situations.

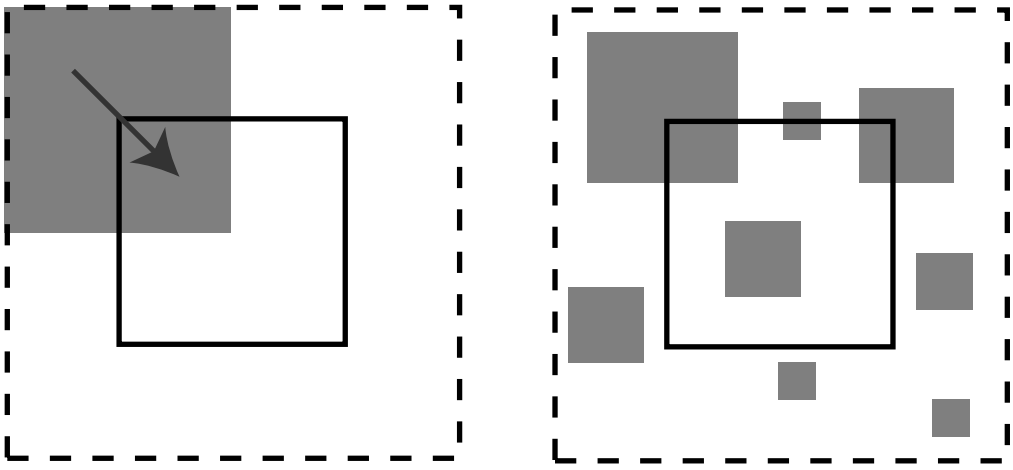

Figure 18. Schematic to show two different arrangements of areas of highest rainfall accumulations within a region. The square enclosed by the continuous line represents a river catchment area. The square enclosed by the dashed line represents an extended area around the catchment. The grey squares represent the highest accumulations within the dashed square that cover the same area as the small square. The arrow shows how the grey area could be translated into the small square in the left hand picture.

Method

1. Steps 1-3 as before
2. Search through all the points in both the catchment area and extended area and find the highest N values of precipitation accumulation.
3. Use a cell detection algorithm to find all the cells within the catchment area and extended area that have points with accumulations within the range of the N highest accumulations.
4. Arrange the cells in order of total accumulation, highest first.
5. Calculate an alternative maximum average catchment accumulation by weighting the contribution of each cell – most weight to the highest total accumulation and dividing by N. The point of this is to make sure that it is not possible to count all the contribution from every cell, especially when there are a large number of cells.
6. Display this value on the catchment area map (or other form of presentation). Repeat 1-5 for all catchments.

The use of the cell-detecting algorithm does slow down the computation, but not restrictively.

Examples

Examples of using the cell weighted maximum average catchment accumulation for extended areas of width 20 and 40km are shown in Figure 19. In this case the weighting given to the cells was 1.0 for the cell with the highest total accumulation, 0.8 for the next highest, 0.8*0.8 for the next, 0.8*0.8*0.8 for the next and so on. A comparison between Figure 19 and Figure 20 shows that the maximum average accumulations have been noticeably reduced for some catchments, but not for most. That is what would be expected for a case like this with a single band of high rainfall totals. Only a few cells are found in each catchment and the impact of a model displacement error could therefore be very significant.
A comparison of this case with the 3rd May 2002 scattered convection event highlights the difference between two different meteorological situations. Figure 20 shows the percentage decrease in the maximum average accumulation when the cell weighting is applied. The percentage change is much greater for the scattered case with most catchments showing a >10-15% change compared to a <5% change for the 11-12th Oct 2000 case in the areas that matter.

The weighting function has been chosen because it seems reasonable and works well, but a more rigorous approach to discovering what this function should look like, should be possible.

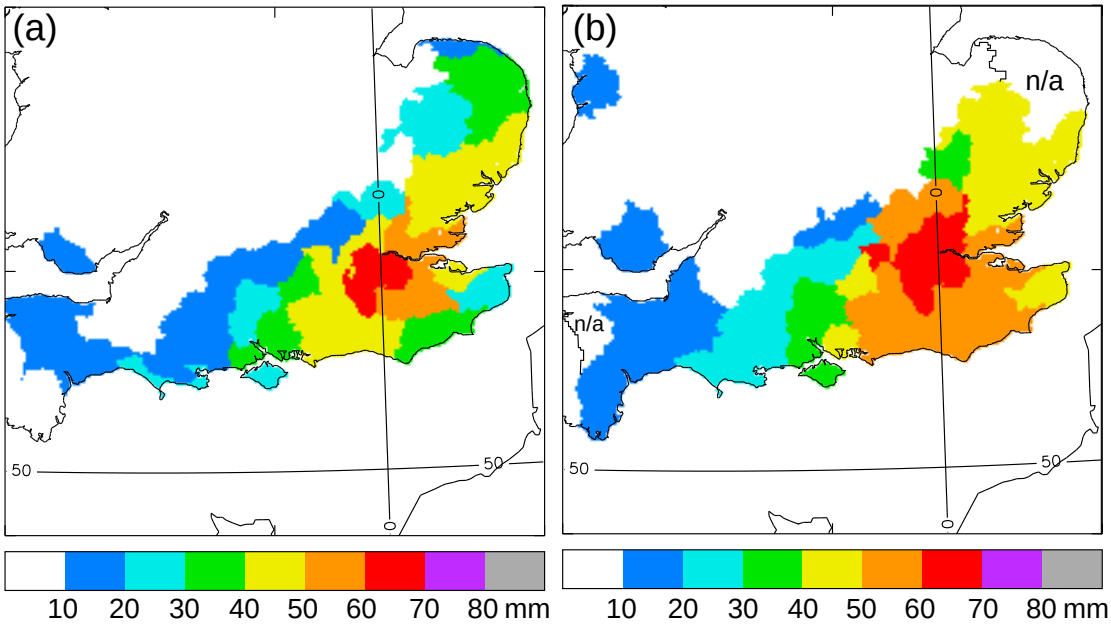

Figure 19. Maximum average catchment rainfall accumulations over an 18-hour period starting at 18UTC 11/10/00 from a 2-km gridlength forecast starting at 18UTC 11/10/00. Accumulations were computed using (a) a 20km extended area with the cell detection switched on and (b) same as (a) but with a 40km extended area. See text for more information.

2. Maximum average accumulations over half catchment areas

The same extended area sampling procedure but this time averaging over the maximum accumulations that comprise half the number of points in a catchment. The values are displayed over the whole catchment. This will give higher values because the area is smaller and is also intended to be used as an indication of the potential for a serious event. This can be computed with or without cell weighting.

Figure 21 shows an example of this for the 20 and 40km extended areas. The catchment highlighted earlier has a maximum half catchment average of >60mm.

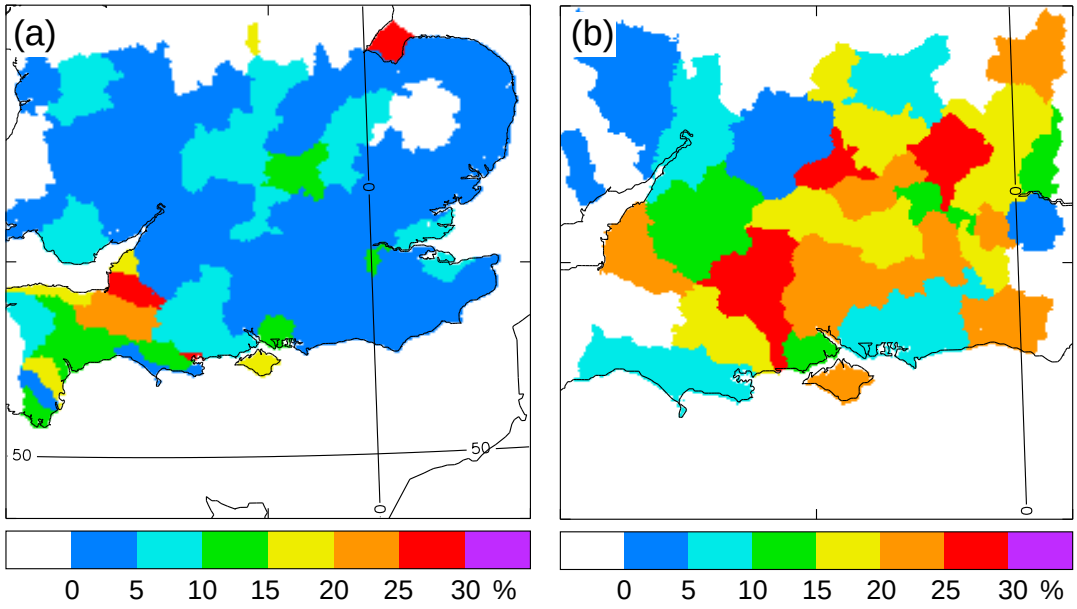

Figure 20. The percentage reduction in the maximum average catchment accumulation when applying the cell-detection algorithm to; (a) A continuous band of rainfall accumulations over a 18-hour period starting at 18UTC 11/10/00 from a 2-km gridlength forecast starting at that time, (b) Scattered rainfall accumulations on 03/05/02, a 9-hour period starting at 09UTC from a 1km gridlength forecast that started at 01UTC.

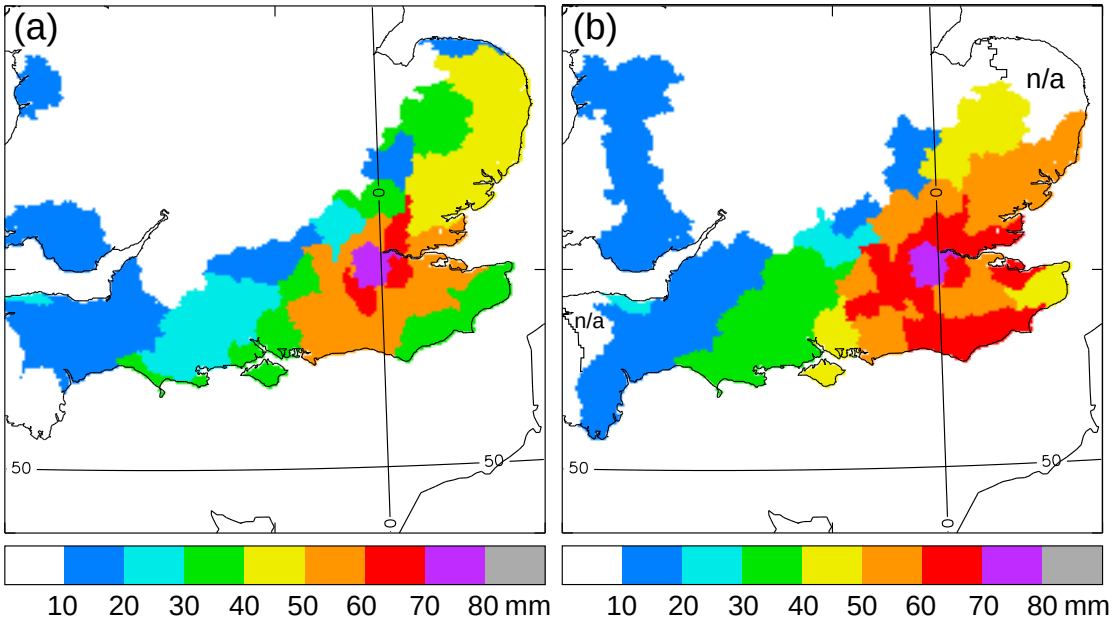

Figure 21. Maximum average half-catchment rainfall accumulations over an 18-hour period starting at 18UTC 11/10/00 from a 2-km gridlength forecast starting at 18UTC 11/10/00. Accumulations were computed using (a) a 20km extended area and (b) a 40km extended area. Accumulations are presented on whole catchments. See text for more information.

8.3.3 Probability of a catchment average accumulation exceeding a threshold

In the same way that point probabilities were computed by finding the proportion of surrounding points exceeding a given threshold, the probability of a catchment average accumulation exceeding a given threshold can be found by finding the proportion of surrounding areas of the same size as the catchment that exceed the threshold.

Method

1. Find the precipitation accumulations over a specified time period at every gridpoint in the forecast domain
2. Locate the gridpoints in a catchment area and define a set of gridpoints surrounding the catchment that are within some pre-determined distance from the edge of the catchment. (As for catchment products in 7.3)
3. Count the number of gridpoints N within the catchment area alone and find the nearest integer to the square root of N to find the length of a side of the square area A that is closest in area to the catchment.
4. For every point within the catchment area and the extended area, calculate the average accumulation over a square area A centered at that point.
5. Calculate the proportion of the squares that have an average accumulation over a specified threshold.

The threshold can be a soft threshold as described in section 7.1.1

Purpose

To give an indication of the chance of a critical amount of rain falling on a catchment.

Parameters

1. Size of the extended area – As for the other catchment products, the size of the extended area should have a width of somewhere between 15 and 60km. The choice should be largely determined by the skill of the forecast system with different widths and the use that will be made of the probabilities. In most circumstances, a larger extended area will produce smaller probabilities for likely events, but provide more warning of unlikely events.

2. Threshold – The threshold should be determined by the user. It would be possible to have a different threshold for every catchment and vary that according to current hydrological conditions. This might mean lowering the threshold for a catchment that is already saturated.

3. Accumulation period – The accumulation period should depend on the user requirements (and the forecast run length).

Examples

Figure 22 shows examples of catchment average accumulation probabilities. The differences between (a) and (b) and between (c) and (d) show the effect of increasing the width of the extended area for this event. The change from an extended width of 20km to 40km has decreased the highest probabilities and increased the lowest. As the threshold is increased from 20mm (a) and (b) to 40mm (c) and (d) the probabilities have become smaller.

Extensions

1. Combine probabilities from different forecasts with different weights applied to each forecast.
2. Average over circles instead of squares, or even masks of the catchment shapes. These extensions would probably yield little extra benefit for the extra cost.
3. Display the accumulation that is exceeded for a given probability. The catchments are then coloured according to the accumulation that is exceeded – each map is for a different probability. A four-up picture might be for probabilities of 0.8, 0.4, 0.2 and 0.1.

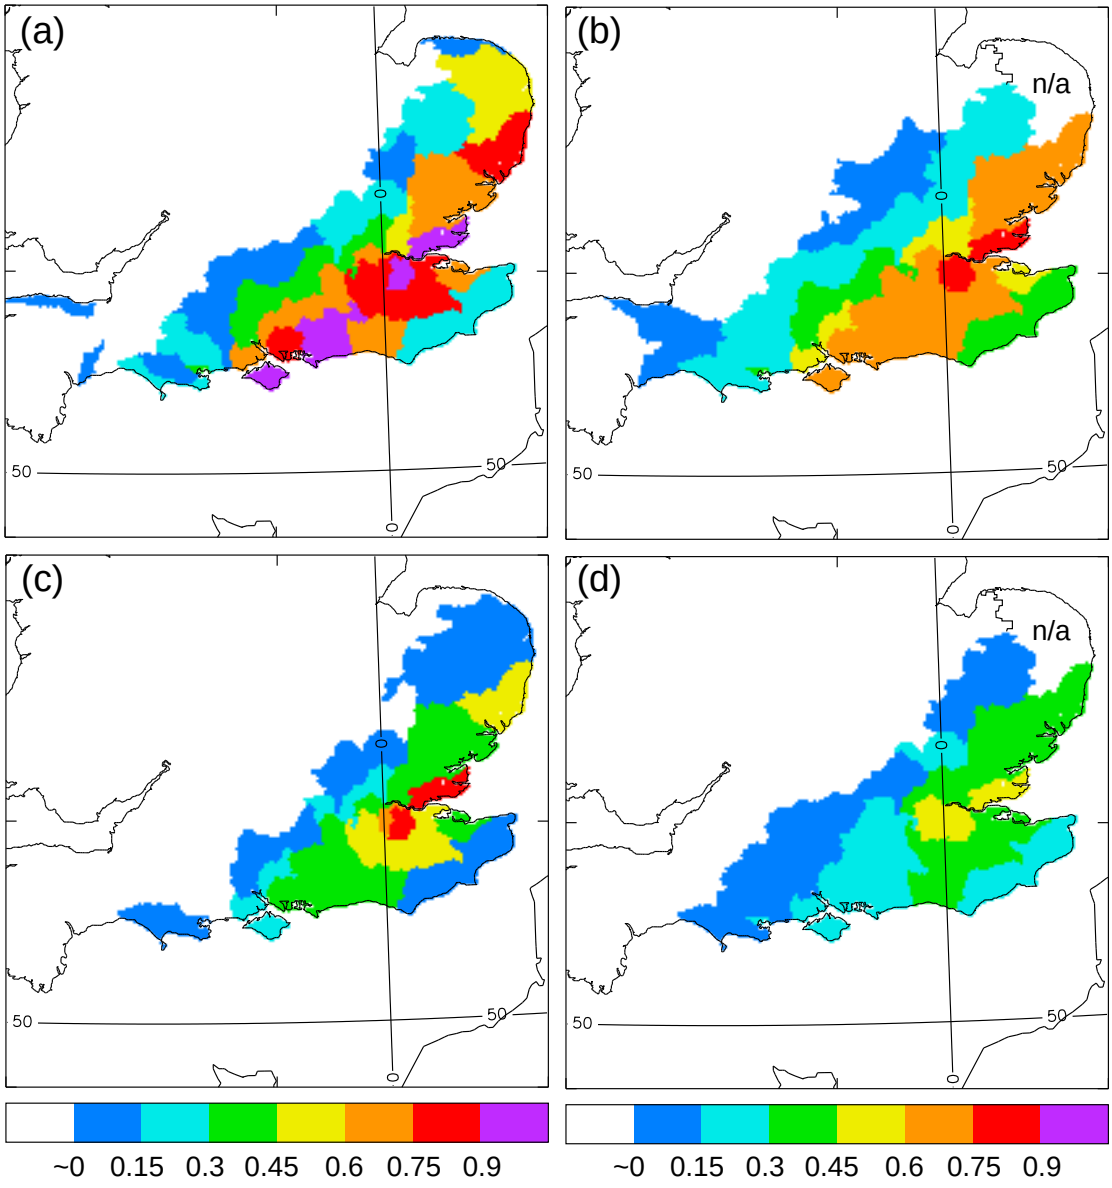

Figure 22. Probabilities of rainfall accumulations exceeding 20 or 40mm over an 18-hour period starting at 18UTC 11/10/00 from a 2-km gridlength forecast starting at 18UTC 11/10/00. Probabilities are for (a) accumulations exceeding 20mm using a 20km extended area; (b) Same as (a) but using a 40km extended area; (c) accumulations exceeding 40mm using a 20km extended area; (d) same as (c) but using a 40km extended area. See text for more information.

8.4 Graphical representation of catchment products

The catchment area precipitation accumulation products can be computed over hourly periods and displayed graphically. Such information could also be fed directly into a hydrological model.

Figure 23 shows an example of the hourly forecast accumulations over the catchment highlighted in Figure 16(a) for the 11-12th Oct 2000 event. Also shown are the hourly maximum average accumulations computed using the method in section 7.3.2. These values could be fed into a hydrological model as an 'expected' and 'worst case' scenario.

Figure 24 presents the same hourly accumulations displayed as cumulative accumulations over the forecast period. Some care needs to be taken in the interpretation of these maximum average accumulations. The final (12UTC) maximum average accumulation here is not the same as would be generated by calculating the maximum average accumulation directly from the 18-hour accumulations over the same extended area. By performing the calculation every hour, rain areas that are moving through the catchment area will be counted more than once, unless they cross the extended area within an hour. The way to make the two more equivalent is to reduce the size of the extended area when calculating every hour. The problem then is that the extended catchment may become too small to account for model spatial errors if the majority of the rain never enters the small area. A compromise can be reached in which the size of the extended catchment for the hourly calculations is set to give a final accumulation which is close to that given by calculating over the full time period with an optimal extended area. The alternative is to calculate the average accumulations and maximum average accumulations over the first hour, then the first two hours, then the first three hours etc, and subtract to get the individual hourly amounts. Half catchment size maximum average accumulations could also be plotted.

Probabilities can be incorporated into the graph of cumulative accumulations. In Figure 24 the probabilities of catchment average 18-hour accumulations exceeding 10mm interval thresholds are displayed. There is a 50% chance of the catchment average accumulation exceeding 30mm. The probability of the forecast accumulation is close to 50% in this example, but need not be for different events. The maximum average accumulation should always have a low probability. Probabilities for accumulations exceeded over shorter time periods could also be presented e.g. every 3 hours.

Figure 25 presents information on a 'box and whisker' diagram for another nearby catchment area. This is one of the presentation methods used by the European Centre for Medium-range Weather Forecasts (ECMWF) for their precipitation forecasts. The boxes show the 25-75-percentile range of accumulations that are exceeded; i.e. the 25% value is the accumulation that is not exceeded by 25% of the catchment-sized sampling squares. The lines show the range of the maximum and minimum accumulations.

It is also possible to plot lines of average accumulation exceeded on a graph of time period against probability, or plot lines of probability on a graph of time period against accumulation exceeded.

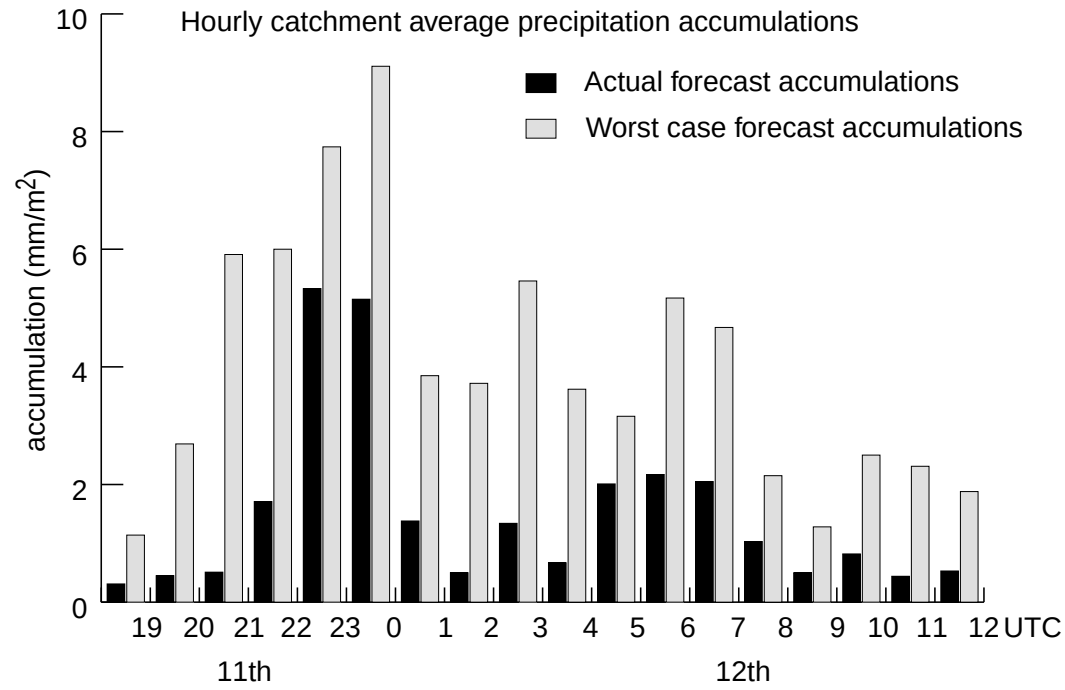

Figure 23. Graph of the hourly average rainfall accumulation over the catchment shown in Figure 16(a) (black bars) and the hourly maximum average rainfall accumulations for the same catchment (grey bars) from a 2-km gridlength model simulation starting at 18UTC 11/10/00.

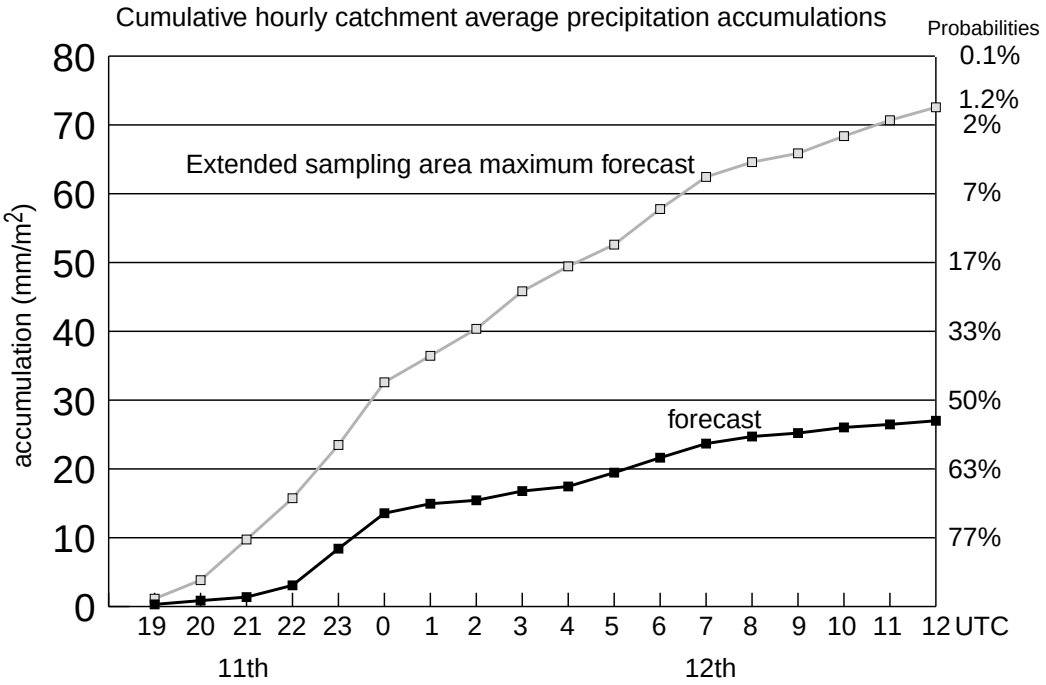

Figure 24. Graph of the cumulative catchment average rainfall accumulations taken from the hourly values in Figure 23. The probabilities are for exceeding 18-hour accumulations thresholds.

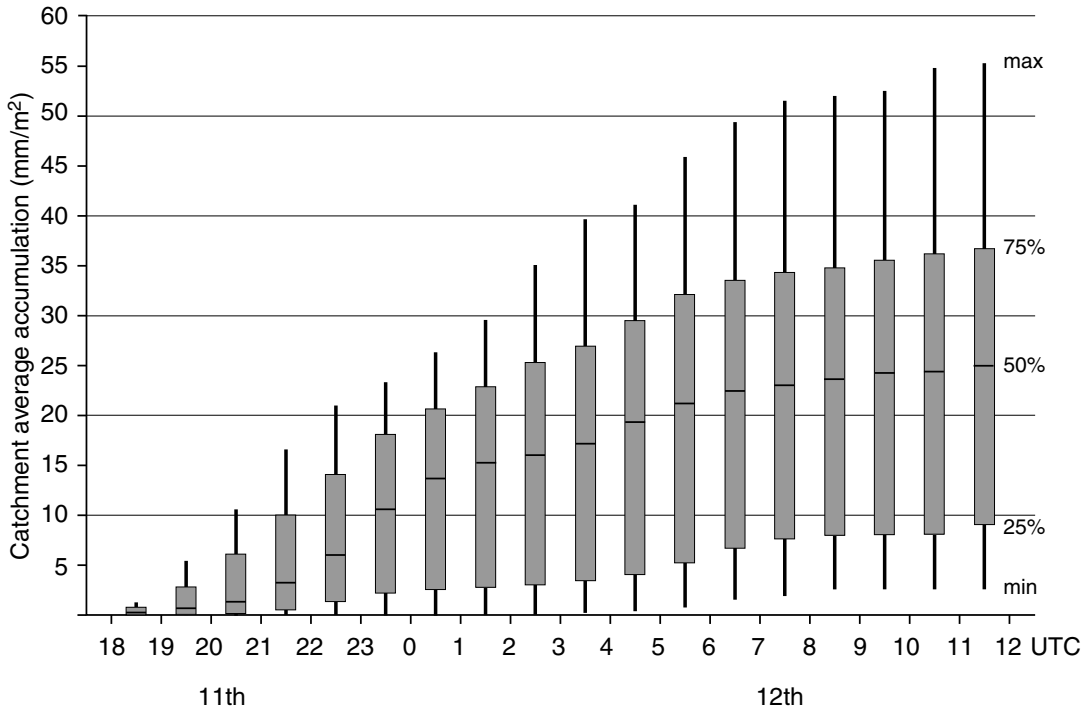

Figure 25. Graph of the cumulative catchment-average rainfall accumulations for a different catchment to that in Figure 24. See text for more information.

9 Final comments

This report has described reasons why diagnostic products are required to present output from high-resolution model precipitation forecasts. Several potentially useful products have been described and examples shown. They are only a small set of what could be developed. It is possible to think of others with an increasing degree of sophistication. For example, a cell tracking as well as locating algorithm would allow the computation of expected rainfall duration within regions or the probability of a period of rain lasting longer than say 30 minutes. Output from 4km gridlength forecasts of rain from the convection scheme could be combined with 1km gridlength forecasts of rain from the model dynamics to produce a wider area of risk of convection alongside the scale and intensity of the convection should it occur. Different types of rainfall event could be classified and used in climatologies. Weighting with distance could be made more formal by the use of a pre-determined masking function based on covariance statistics as an alternative to the combined averaging or a recursive filter.

For now though, it is probably better to examine what is already available. The products described in this report need to be incorporated into the verification of a large number of high-resolution forecasts so that they can be tuned to the most appropriate scales. They should be presented to potential customers to hear their views on how useful they might be and what might be better. A set of products could then be included into a storm-scale modelling system, and tailored to individual customers' needs.

The choice of averaging areas presented in this report has been based on subjective judgement, alongside a trial and error approach to testing the behaviour of the diagnostics. Further work is needed to develop a verification methodology that will make this decision more objective.

10 References

Forbes, R. M. and Clark, P. A., 2003: Sensitivity of extratropical cyclone mesoscale structure to the parametrization of ice microphysical processes., *Q.J.R.Meteorol.Soc,* **129**, 1123-1148.

Golding, B.W. 1998: Nimrod: A system for generating automated very short range forecasts., Meteorol. Appl. 5, pp1-16

Hand, W.H. 2002: The Met Office Convection Diagnosis Scheme., Meteorol. Appl. 9, pp69-83

Lean, H. W., 2003: High Resolution Trial Model project, results following the completion of stage 3. Available on request.

Roberts, N. M., 2003a: Results from high-resolution simulations of convective events. *JCMM internal report.,* 140. *(NWP Technical Report.,* 402*)*

Roberts, N. M., 2003b: The impact of a change to the use of the convection scheme in high-resolution simulations of convective events. *JCMM internal report.,* 142. *(NWP Technical Report.,* 407*)*

Zhang, F., Snyder, S. and Rotunno, R., 2003: Effects of Moist Convection on Mesoscale Predictability., *J. Atmos. Sci.*, 60, 1173-1185

11 Acknowledgements

This work was partially funded by Defra under the Flood Forecasting and Warning Theme. I would particularly like to thank Brian Golding and Peter for extremely useful comments. Thanks to David Stephenson for discussion about the use of odds as an alternative to probabilities. Thanks to Adrian Thompkins for the use of his cell-detection code.